



Cisco Security Appliance Command Line Configuration Guide

For the Cisco ASA 5500 Series and Cisco PIX 500 Series

Software Version 7.2

Americas Headquarters

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

Customer Order Number: N/A, Online only
Text Part Number: OL-10088-02

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

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

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

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

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

CCDE, CCSI, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0903R)

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

Cisco Security Appliance Command Line Configuration Guide
Copyright © 2008 Cisco Systems, Inc. All rights reserved.



CONTENTS

About This Guide xxxv

Document Objectives xxxv

Audience xxxv

Related Documentation xxxvi

Document Organization xxxvi

Document Conventions xxxix

Obtaining Documentation and Submitting a Service Request xxxix

1-xi

PART 1

Getting Started and General Information

CHAPTER 1

Introduction to the Security Appliance 1-1

Firewall Functional Overview 1-1

Security Policy Overview 1-2

Permitting or Denying Traffic with Access Lists 1-2

Applying NAT 1-2

Using AAA for Through Traffic 1-2

Applying HTTP, HTTPS, or FTP Filtering 1-3

Applying Application Inspection 1-3

Sending Traffic to the Advanced Inspection and Prevention Security Services Module 1-3

Sending Traffic to the Content Security and Control Security Services Module 1-3

Applying QoS Policies 1-3

Applying Connection Limits and TCP Normalization 1-3

Firewall Mode Overview 1-3

Stateful Inspection Overview 1-4

VPN Functional Overview 1-5

Intrusion Prevention Services Functional Overview 1-5

Security Context Overview 1-6

CHAPTER 2

Getting Started 2-1

Getting Started with Your Platform Model 2-1

Factory Default Configurations 2-1

Restoring the Factory Default Configuration 2-2

ASA 5505 Default Configuration	2-2
ASA 5510 and Higher Default Configuration	2-3
PIX 515/515E Default Configuration	2-4
Accessing the Command-Line Interface	2-4
Setting Transparent or Routed Firewall Mode	2-5
Working with the Configuration	2-6
Saving Configuration Changes	2-6
Saving Configuration Changes in Single Context Mode	2-7
Saving Configuration Changes in Multiple Context Mode	2-7
Copying the Startup Configuration to the Running Configuration	2-8
Viewing the Configuration	2-8
Clearing and Removing Configuration Settings	2-9
Creating Text Configuration Files Offline	2-9

CHAPTER 3

Enabling Multiple Context Mode	3-1
Security Context Overview	3-1
Common Uses for Security Contexts	3-1
Unsupported Features	3-2
Context Configuration Files	3-2
Context Configurations	3-2
System Configuration	3-2
Admin Context Configuration	3-2
How the Security Appliance Classifies Packets	3-3
Valid Classifier Criteria	3-3
Invalid Classifier Criteria	3-4
Classification Examples	3-5
Cascading Security Contexts	3-8
Management Access to Security Contexts	3-9
System Administrator Access	3-9
Context Administrator Access	3-10
Enabling or Disabling Multiple Context Mode	3-10
Backing Up the Single Mode Configuration	3-10
Enabling Multiple Context Mode	3-10
Restoring Single Context Mode	3-11

CHAPTER 4

Configuring Switch Ports and VLAN Interfaces for the Cisco ASA 5505 Adaptive Security Appliance	4-1
Interface Overview	4-1
Understanding ASA 5505 Ports and Interfaces	4-2

Maximum Active VLAN Interfaces for Your License	4-2
Default Interface Configuration	4-4
VLAN MAC Addresses	4-4
Power Over Ethernet	4-4
Monitoring Traffic Using SPAN	4-4
Security Level Overview	4-5
Configuring VLAN Interfaces	4-5
Configuring Switch Ports as Access Ports	4-9
Configuring a Switch Port as a Trunk Port	4-11
Allowing Communication Between VLAN Interfaces on the Same Security Level	4-13

CHAPTER 5**Configuring Ethernet Settings and Subinterfaces 5-1**

Configuring and Enabling RJ-45 Interfaces	5-1
Configuring and Enabling Fiber Interfaces	5-3
Configuring and Enabling VLAN Subinterfaces and 802.1Q Trunking	5-3

CHAPTER 6**Adding and Managing Security Contexts 6-1**

Configuring Resource Management	6-1
Classes and Class Members Overview	6-1
Resource Limits	6-2
Default Class	6-3
Class Members	6-4
Configuring a Class	6-4
Configuring a Security Context	6-7
Automatically Assigning MAC Addresses to Context Interfaces	6-11
Changing Between Contexts and the System Execution Space	6-11
Managing Security Contexts	6-12
Removing a Security Context	6-12
Changing the Admin Context	6-13
Changing the Security Context URL	6-13
Reloading a Security Context	6-14
Reloading by Clearing the Configuration	6-14
Reloading by Removing and Re-adding the Context	6-15
Monitoring Security Contexts	6-15
Viewing Context Information	6-15
Viewing Resource Allocation	6-16
Viewing Resource Usage	6-19
Monitoring SYN Attacks in Contexts	6-20

CHAPTER 7

Configuring Interface Parameters 7-1

- Security Level Overview 7-1
- Configuring the Interface 7-2
- Allowing Communication Between Interfaces on the Same Security Level 7-6

CHAPTER 8

Configuring Basic Settings 8-1

- Changing the Login Password 8-1
- Changing the Enable Password 8-1
- Setting the Hostname 8-2
- Setting the Domain Name 8-2
- Setting the Date and Time 8-2
 - Setting the Time Zone and Daylight Saving Time Date Range 8-3
 - Setting the Date and Time Using an NTP Server 8-4
 - Setting the Date and Time Manually 8-5
- Setting the Management IP Address for a Transparent Firewall 8-5

CHAPTER 9

Configuring IP Routing 9-1

- How Routing Behaves Within the ASA Security Appliance 9-1
 - Egress Interface Selection Process 9-1
 - Next Hop Selection Process 9-2
- Configuring Static and Default Routes 9-2
 - Configuring a Static Route 9-3
 - Configuring a Default Route 9-4
 - Configuring Static Route Tracking 9-5
- Defining Route Maps 9-7
- Configuring OSPF 9-8
 - OSPF Overview 9-9
 - Enabling OSPF 9-10
 - Redistributing Routes Into OSPF 9-10
 - Configuring OSPF Interface Parameters 9-11
 - Configuring OSPF Area Parameters 9-13
 - Configuring OSPF NSSA 9-14
 - Configuring Route Summarization Between OSPF Areas 9-15
 - Configuring Route Summarization When Redistributing Routes into OSPF 9-16
 - Defining Static OSPF Neighbors 9-16
 - Generating a Default Route 9-17
 - Configuring Route Calculation Timers 9-17
 - Logging Neighbors Going Up or Down 9-18

Displaying OSPF Update Packet Pacing	9-19
Monitoring OSPF	9-19
Restarting the OSPF Process	9-20
Configuring RIP	9-20
Enabling and Configuring RIP	9-20
Redistributing Routes into the RIP Routing Process	9-22
Configuring RIP Send/Receive Version on an Interface	9-22
Enabling RIP Authentication	9-23
Monitoring RIP	9-23
The Routing Table	9-24
Displaying the Routing Table	9-24
How the Routing Table is Populated	9-24
Backup Routes	9-26
How Forwarding Decisions are Made	9-26
Dynamic Routing and Failover	9-26

CHAPTER 10
Configuring DHCP, DDNS, and WCCP Services 10-1

Configuring a DHCP Server	10-1
Enabling the DHCP Server	10-2
Configuring DHCP Options	10-3
Using Cisco IP Phones with a DHCP Server	10-4
Configuring DHCP Relay Services	10-5
Configuring Dynamic DNS	10-6
Example 1: Client Updates Both A and PTR RRs for Static IP Addresses	10-7
Example 2: Client Updates Both A and PTR RRs; DHCP Server Honors Client Update Request; FQDN Provided Through Configuration	10-7
Example 3: Client Includes FQDN Option Instructing Server Not to Update Either RR; Server Overrides Client and Updates Both RRs.	10-8
Example 4: Client Asks Server To Perform Both Updates; Server Configured to Update PTR RR Only; Honors Client Request and Updates Both A and PTR RR	10-8
Example 5: Client Updates A RR; Server Updates PTR RR	10-9
Configuring Web Cache Services Using WCCP	10-9
WCCP Feature Support	10-9
WCCP Interaction With Other Features	10-10
Enabling WCCP Redirection	10-10

CHAPTER 11
Configuring Multicast Routing 11-13

Multicast Routing Overview	11-13
Enabling Multicast Routing	11-14

Configuring IGMP Features	11-14
Disabling IGMP on an Interface	11-15
Configuring Group Membership	11-15
Configuring a Statically Joined Group	11-15
Controlling Access to Multicast Groups	11-15
Limiting the Number of IGMP States on an Interface	11-16
Modifying the Query Interval and Query Timeout	11-16
Changing the Query Response Time	11-17
Changing the IGMP Version	11-17
Configuring Stub Multicast Routing	11-17
Configuring a Static Multicast Route	11-17
Configuring PIM Features	11-18
Disabling PIM on an Interface	11-18
Configuring a Static Rendezvous Point Address	11-19
Configuring the Designated Router Priority	11-19
Filtering PIM Register Messages	11-19
Configuring PIM Message Intervals	11-20
Configuring a Multicast Boundary	11-20
Filtering PIM Neighbors	11-20
Supporting Mixed Bidirectional/Sparse-Mode PIM Networks	11-21
For More Information about Multicast Routing	11-22

CHAPTER 12

Configuring IPv6 12-1

IPv6-enabled Commands	12-1
Configuring IPv6	12-2
Configuring IPv6 on an Interface	12-3
Configuring a Dual IP Stack on an Interface	12-4
Enforcing the Use of Modified EUI-64 Interface IDs in IPv6 Addresses	12-4
Configuring IPv6 Duplicate Address Detection	12-4
Configuring IPv6 Default and Static Routes	12-5
Configuring IPv6 Access Lists	12-6
Configuring IPv6 Neighbor Discovery	12-7
Configuring Neighbor Solicitation Messages	12-7
Configuring Router Advertisement Messages	12-9
Multicast Listener Discovery Support	12-11
Configuring a Static IPv6 Neighbor	12-11
Verifying the IPv6 Configuration	12-11
The show ipv6 interface Command	12-12
The show ipv6 route Command	12-12

The show ipv6 mld traffic Command 12-13

CHAPTER 13

Configuring AAA Servers and the Local Database 13-1

AAA Overview 13-1

About Authentication 13-1

About Authorization 13-2

About Accounting 13-2

AAA Server and Local Database Support 13-2

Summary of Support 13-3

RADIUS Server Support 13-3

Authentication Methods 13-4

Attribute Support 13-4

RADIUS Authorization Functions 13-4

TACACS+ Server Support 13-4

SDI Server Support 13-4

SDI Version Support 13-5

Two-step Authentication Process 13-5

SDI Primary and Replica Servers 13-5

NT Server Support 13-5

Kerberos Server Support 13-5

LDAP Server Support 13-6

Authentication with LDAP 13-6

Authorization with LDAP for VPN 13-7

LDAP Attribute Mapping 13-8

SSO Support for WebVPN with HTTP Forms 13-9

Local Database Support 13-9

User Profiles 13-10

Fallback Support 13-10

Configuring the Local Database 13-10

Identifying AAA Server Groups and Servers 13-12

Using Certificates and User Login Credentials 13-15

Using User Login Credentials 13-15

Using certificates 13-16

Supporting a Zone Labs Integrity Server 13-16

Overview of Integrity Server and Security Appliance Interaction 13-17

Configuring Integrity Server Support 13-17

CHAPTER 14

Configuring Failover 14-1

Understanding Failover 14-1

Failover System Requirements	14-2
Hardware Requirements	14-2
Software Requirements	14-2
License Requirements	14-2
The Failover and Stateful Failover Links	14-3
Failover Link	14-3
Stateful Failover Link	14-5
Active/Active and Active/Standby Failover	14-6
Active/Standby Failover	14-6
Active/Active Failover	14-10
Determining Which Type of Failover to Use	14-15
Regular and Stateful Failover	14-15
Regular Failover	14-16
Stateful Failover	14-16
Failover Health Monitoring	14-16
Unit Health Monitoring	14-17
Interface Monitoring	14-17
Failover Feature/Platform Matrix	14-18
Failover Times by Platform	14-18
Configuring Failover	14-19
Failover Configuration Limitations	14-19
Configuring Active/Standby Failover	14-19
Prerequisites	14-20
Configuring Cable-Based Active/Standby Failover (PIX Security Appliance Only)	14-20
Configuring LAN-Based Active/Standby Failover	14-21
Configuring Optional Active/Standby Failover Settings	14-25
Configuring Active/Active Failover	14-27
Prerequisites	14-27
Configuring Cable-Based Active/Active Failover (PIX security appliance)	14-27
Configuring LAN-Based Active/Active Failover	14-29
Configuring Optional Active/Active Failover Settings	14-33
Configuring Unit Health Monitoring	14-39
Configuring Failover Communication Authentication/Encryption	14-39
Verifying the Failover Configuration	14-40
Using the show failover Command	14-40
Viewing Monitored Interfaces	14-48
Displaying the Failover Commands in the Running Configuration	14-48
Testing the Failover Functionality	14-49
Controlling and Monitoring Failover	14-49
Forcing Failover	14-49

Disabling Failover	14-50
Restoring a Failed Unit or Failover Group	14-50
Monitoring Failover	14-50
Failover System Messages	14-51
Debug Messages	14-51
SNMP	14-51

PART 2**Configuring the Firewall****CHAPTER 15**

Firewall Mode Overview	15-1
Routed Mode Overview	15-1
IP Routing Support	15-1
Network Address Translation	15-2
How Data Moves Through the Security Appliance in Routed Firewall Mode	15-3
An Inside User Visits a Web Server	15-3
An Outside User Visits a Web Server on the DMZ	15-4
An Inside User Visits a Web Server on the DMZ	15-6
An Outside User Attempts to Access an Inside Host	15-7
A DMZ User Attempts to Access an Inside Host	15-8
Transparent Mode Overview	15-8
Transparent Firewall Network	15-9
Allowing Layer 3 Traffic	15-9
Allowed MAC Addresses	15-9
Passing Traffic Not Allowed in Routed Mode	15-9
MAC Address Lookups	15-10
Using the Transparent Firewall in Your Network	15-10
Transparent Firewall Guidelines	15-10
Unsupported Features in Transparent Mode	15-11
How Data Moves Through the Transparent Firewall	15-13
An Inside User Visits a Web Server	15-14
An Outside User Visits a Web Server on the Inside Network	15-15
An Outside User Attempts to Access an Inside Host	15-16

CHAPTER 16

Identifying Traffic with Access Lists	16-1
Access List Overview	16-1
Access List Types	16-2
Access Control Entry Order	16-2
Access Control Implicit Deny	16-3
IP Addresses Used for Access Lists When You Use NAT	16-3

Adding an Extended Access List	16-5
Extended Access List Overview	16-5
Allowing Broadcast and Multicast Traffic through the Transparent Firewall	16-6
Adding an Extended ACE	16-6
Adding an EtherType Access List	16-8
EtherType Access List Overview	16-8
Supported EtherTypes	16-8
Implicit Permit of IP and ARPs Only	16-9
Implicit and Explicit Deny ACE at the End of an Access List	16-9
IPv6 Unsupported	16-9
Using Extended and EtherType Access Lists on the Same Interface	16-9
Allowing MPLS	16-9
Adding an EtherType ACE	16-10
Adding a Standard Access List	16-11
Adding a Webtype Access List	16-11
Simplifying Access Lists with Object Grouping	16-11
How Object Grouping Works	16-12
Adding Object Groups	16-12
Adding a Protocol Object Group	16-13
Adding a Network Object Group	16-13
Adding a Service Object Group	16-14
Adding an ICMP Type Object Group	16-15
Nesting Object Groups	16-15
Using Object Groups with an Access List	16-16
Displaying Object Groups	16-17
Removing Object Groups	16-17
Adding Remarks to Access Lists	16-18
Scheduling Extended Access List Activation	16-18
Adding a Time Range	16-18
Applying the Time Range to an ACE	16-19
Logging Access List Activity	16-20
Access List Logging Overview	16-20
Configuring Logging for an Access Control Entry	16-21
Managing Deny Flows	16-22

CHAPTER 17

Applying NAT 17-1

NAT Overview	17-1
Introduction to NAT	17-2
NAT Control	17-3

NAT Types	17-5
Dynamic NAT	17-5
PAT	17-7
Static NAT	17-7
Static PAT	17-8
Bypassing NAT When NAT Control is Enabled	17-9
Policy NAT	17-9
NAT and Same Security Level Interfaces	17-13
Order of NAT Commands Used to Match Real Addresses	17-14
Mapped Address Guidelines	17-14
DNS and NAT	17-14
Configuring NAT Control	17-16
Using Dynamic NAT and PAT	17-17
Dynamic NAT and PAT Implementation	17-17
Configuring Dynamic NAT or PAT	17-23
Using Static NAT	17-26
Using Static PAT	17-27
Bypassing NAT	17-29
Configuring Identity NAT	17-30
Configuring Static Identity NAT	17-30
Configuring NAT Exemption	17-32
NAT Examples	17-33
Overlapping Networks	17-34
Redirecting Ports	17-35

CHAPTER 18
Permitting or Denying Network Access 18-1

Inbound and Outbound Access List Overview	18-1
Applying an Access List to an Interface	18-2

CHAPTER 19
Applying AAA for Network Access 19-1

AAA Performance	19-1
Configuring Authentication for Network Access	19-1
Authentication Overview	19-2
One-Time Authentication	19-2
Applications Required to Receive an Authentication Challenge	19-2
Security Appliance Authentication Prompts	19-2
Static PAT and HTTP	19-3
Enabling Network Access Authentication	19-3

Enabling Secure Authentication of Web Clients	19-5
Authenticating Directly with the Security Appliance	19-6
Enabling Direct Authentication Using HTTP and HTTPS	19-6
Enabling Direct Authentication Using Telnet	19-6
Configuring Authorization for Network Access	19-6
Configuring TACACS+ Authorization	19-7
Configuring RADIUS Authorization	19-8
Configuring a RADIUS Server to Send Downloadable Access Control Lists	19-9
Configuring a RADIUS Server to Download Per-User Access Control List Names	19-12
Configuring Accounting for Network Access	19-13
Using MAC Addresses to Exempt Traffic from Authentication and Authorization	19-14

CHAPTER 20**Applying Filtering Services 20-1**

Filtering Overview	20-1
Filtering ActiveX Objects	20-2
ActiveX Filtering Overview	20-2
Enabling ActiveX Filtering	20-2
Filtering Java Applets	20-3
Filtering URLs and FTP Requests with an External Server	20-4
URL Filtering Overview	20-4
Identifying the Filtering Server	20-4
Buffering the Content Server Response	20-6
Caching Server Addresses	20-6
Filtering HTTP URLs	20-7
Configuring HTTP Filtering	20-7
Enabling Filtering of Long HTTP URLs	20-7
Truncating Long HTTP URLs	20-7
Exempting Traffic from Filtering	20-8
Filtering HTTPS URLs	20-8
Filtering FTP Requests	20-9
Viewing Filtering Statistics and Configuration	20-9
Viewing Filtering Server Statistics	20-10
Viewing Buffer Configuration and Statistics	20-11
Viewing Caching Statistics	20-11
Viewing Filtering Performance Statistics	20-11
Viewing Filtering Configuration	20-12

CHAPTER 21**Using Modular Policy Framework 21-1**

- Modular Policy Framework Overview 21-1
 - Modular Policy Framework Features 21-1
 - Modular Policy Framework Configuration Overview 21-2
 - Default Global Policy 21-3
- Identifying Traffic (Layer 3/4 Class Map) 21-4
 - Default Class Maps 21-4
 - Creating a Layer 3/4 Class Map for Through Traffic 21-5
 - Creating a Layer 3/4 Class Map for Management Traffic 21-7
- Configuring Special Actions for Application Inspections (Inspection Policy Map) 21-7
 - Inspection Policy Map Overview 21-8
 - Defining Actions in an Inspection Policy Map 21-8
 - Identifying Traffic in an Inspection Class Map 21-11
 - Creating a Regular Expression 21-12
 - Creating a Regular Expression Class Map 21-14
- Defining Actions (Layer 3/4 Policy Map) 21-15
 - Layer 3/4 Policy Map Overview 21-15
 - Policy Map Guidelines 21-16
 - Supported Feature Types 21-16
 - Hierarchical Policy Maps 21-16
 - Feature Directionality 21-17
 - Feature Matching Guidelines within a Policy Map 21-17
 - Feature Matching Guidelines for multiple Policy Maps 21-18
 - Order in Which Multiple Feature Actions are Applied 21-18
 - Default Layer 3/4 Policy Map 21-18
 - Adding a Layer 3/4 Policy Map 21-19
- Applying Actions to an Interface (Service Policy) 21-21
- Modular Policy Framework Examples 21-21
 - Applying Inspection and QoS Policing to HTTP Traffic 21-22
 - Applying Inspection to HTTP Traffic Globally 21-22
 - Applying Inspection and Connection Limits to HTTP Traffic to Specific Servers 21-23
 - Applying Inspection to HTTP Traffic with NAT 21-24

CHAPTER 22**Managing AIP SSM and CSC SSM 22-1**

- Managing the AIP SSM 22-1
 - About the AIP SSM 22-1
 - Getting Started with the AIP SSM 22-2
 - Diverting Traffic to the AIP SSM 22-2
 - Sessioning to the AIP SSM and Running Setup 22-4

Managing the CSC SSM	22-5
About the CSC SSM	22-5
Getting Started with the CSC SSM	22-7
Determining What Traffic to Scan	22-9
Limiting Connections Through the CSC SSM	22-11
Diverting Traffic to the CSC SSM	22-11
Checking SSM Status	22-13
Transferring an Image onto an SSM	22-14

CHAPTER 23

Preventing Network Attacks 23-1

Configuring TCP Normalization	23-1
TCP Normalization Overview	23-1
Enabling the TCP Normalizer	23-2
Configuring Connection Limits and Timeouts	23-6
Connection Limit Overview	23-7
TCP Intercept Overview	23-7
Disabling TCP Intercept for Management Packets for Clientless SSL Compatibility	23-7
Dead Connection Detection (DCD) Overview	23-7
TCP Sequence Randomization Overview	23-8
Enabling Connection Limits and Timeouts	23-8
Preventing IP Spoofing	23-10
Configuring the Fragment Size	23-11
Blocking Unwanted Connections	23-11
Configuring IP Audit for Basic IPS Support	23-12

CHAPTER 24

Configuring QoS 24-1

QoS Overview	24-1
Supported QoS Features	24-2
What is a Token Bucket?	24-2
Policing Overview	24-3
Priority Queueing Overview	24-3
Traffic Shaping Overview	24-4
How QoS Features Interact	24-4
DSCP and DiffServ Preservation	24-5
Creating the Standard Priority Queue for an Interface	24-5
Determining the Queue and TX Ring Limits	24-6
Configuring the Priority Queue	24-7
Identifying Traffic for QoS Using Class Maps	24-8

Creating a QoS Class Map	24-8
QoS Class Map Examples	24-8
Creating a Policy for Standard Priority Queueing and/or Policing	24-9
Creating a Policy for Traffic Shaping and Hierarchical Priority Queueing	24-11
Viewing QoS Statistics	24-13
Viewing QoS Police Statistics	24-13
Viewing QoS Standard Priority Statistics	24-14
Viewing QoS Shaping Statistics	24-14
Viewing QoS Standard Priority Queue Statistics	24-15

CHAPTER 25

Configuring Application Layer Protocol Inspection 25-1

Inspection Engine Overview	25-2
When to Use Application Protocol Inspection	25-2
Inspection Limitations	25-2
Default Inspection Policy	25-3
Configuring Application Inspection	25-5
CTIQBE Inspection	25-9
CTIQBE Inspection Overview	25-9
Limitations and Restrictions	25-10
Verifying and Monitoring CTIQBE Inspection	25-10
DCERPC Inspection	25-11
DCERPC Overview	25-11
Configuring a DCERPC Inspection Policy Map for Additional Inspection Control	25-12
DNS Inspection	25-13
How DNS Application Inspection Works	25-13
How DNS Rewrite Works	25-14
Configuring DNS Rewrite	25-15
Using the Static Command for DNS Rewrite	25-15
Using the Alias Command for DNS Rewrite	25-16
Configuring DNS Rewrite with Two NAT Zones	25-16
DNS Rewrite with Three NAT Zones	25-17
Configuring DNS Rewrite with Three NAT Zones	25-19
Verifying and Monitoring DNS Inspection	25-20
Configuring a DNS Inspection Policy Map for Additional Inspection Control	25-20
ESMTP Inspection	25-23
Configuring an ESMTP Inspection Policy Map for Additional Inspection Control	25-24
FTP Inspection	25-26
FTP Inspection Overview	25-27

Using the strict Option	25-27
Configuring an FTP Inspection Policy Map for Additional Inspection Control	25-28
Verifying and Monitoring FTP Inspection	25-31
GTP Inspection	25-32
GTP Inspection Overview	25-32
Configuring a GTP Inspection Policy Map for Additional Inspection Control	25-33
Verifying and Monitoring GTP Inspection	25-37
H.323 Inspection	25-38
H.323 Inspection Overview	25-38
How H.323 Works	25-38
Limitations and Restrictions	25-39
Configuring an H.323 Inspection Policy Map for Additional Inspection Control	25-40
Configuring H.323 and H.225 Timeout Values	25-42
Verifying and Monitoring H.323 Inspection	25-43
Monitoring H.225 Sessions	25-43
Monitoring H.245 Sessions	25-43
Monitoring H.323 RAS Sessions	25-44
HTTP Inspection	25-44
HTTP Inspection Overview	25-44
Configuring an HTTP Inspection Policy Map for Additional Inspection Control	25-45
Instant Messaging Inspection	25-49
IM Inspection Overview	25-49
Configuring an Instant Messaging Inspection Policy Map for Additional Inspection Control	25-49
ICMP Inspection	25-52
ICMP Error Inspection	25-52
ILS Inspection	25-53
IPSec Pass Through Inspection	25-54
IPSec Pass Through Inspection Overview	25-54
Configuring an IPSec Pass Through Inspection Policy Map for Additional Inspection Control	25-54
MGCP Inspection	25-56
MGCP Inspection Overview	25-56
Configuring an MGCP Inspection Policy Map for Additional Inspection Control	25-58
Configuring MGCP Timeout Values	25-59
Verifying and Monitoring MGCP Inspection	25-59
NetBIOS Inspection	25-60
Configuring a NetBIOS Inspection Policy Map for Additional Inspection Control	25-60
PPTP Inspection	25-62
RADIUS Accounting Inspection	25-62

Configuring a RADIUS Inspection Policy Map for Additional Inspection Control	25-63
RSH Inspection	25-63
RTSP Inspection	25-63
RTSP Inspection Overview	25-63
Using RealPlayer	25-64
Restrictions and Limitations	25-64
SIP Inspection	25-65
SIP Inspection Overview	25-65
SIP Instant Messaging	25-65
Configuring a SIP Inspection Policy Map for Additional Inspection Control	25-66
Configuring SIP Timeout Values	25-70
Verifying and Monitoring SIP Inspection	25-70
Skinny (SCCP) Inspection	25-71
SCCP Inspection Overview	25-71
Supporting Cisco IP Phones	25-71
Restrictions and Limitations	25-72
Verifying and Monitoring SCCP Inspection	25-72
Configuring a Skinny (SCCP) Inspection Policy Map for Additional Inspection Control	25-73
SMTP and Extended SMTP Inspection	25-74
SNMP Inspection	25-76
SQL*Net Inspection	25-76
Sun RPC Inspection	25-77
Sun RPC Inspection Overview	25-77
Managing Sun RPC Services	25-77
Verifying and Monitoring Sun RPC Inspection	25-78
TFTP Inspection	25-79
XDMCP Inspection	25-80

CHAPTER 26

Configuring ARP Inspection and Bridging Parameters 26-1

Configuring ARP Inspection	26-1
ARP Inspection Overview	26-1
Adding a Static ARP Entry	26-2
Enabling ARP Inspection	26-2
Customizing the MAC Address Table	26-3
MAC Address Table Overview	26-3
Adding a Static MAC Address	26-3
Setting the MAC Address Timeout	26-4
Disabling MAC Address Learning	26-4

[Viewing the MAC Address Table](#) 26-4

PART 3

Configuring VPN

CHAPTER 27

Configuring IPsec and ISAKMP 27-1

[Tunneling Overview](#) 27-1

[IPsec Overview](#) 27-2

[Configuring ISAKMP](#) 27-2

[ISAKMP Overview](#) 27-2

[Configuring ISAKMP Policies](#) 27-5

[Enabling ISAKMP on the Outside Interface](#) 27-6

[Disabling ISAKMP in Aggressive Mode](#) 27-6

[Determining an ID Method for ISAKMP Peers](#) 27-6

[Enabling IPsec over NAT-T](#) 27-7

[Using NAT-T](#) 27-7

[Enabling IPsec over TCP](#) 27-8

[Waiting for Active Sessions to Terminate Before Rebooting](#) 27-9

[Alerting Peers Before Disconnecting](#) 27-9

[Configuring Certificate Group Matching](#) 27-9

[Creating a Certificate Group Matching Rule and Policy](#) 27-10

[Using the Tunnel-group-map default-group Command](#) 27-11

[Configuring IPsec](#) 27-11

[Understanding IPsec Tunnels](#) 27-11

[Understanding Transform Sets](#) 27-12

[Defining Crypto Maps](#) 27-12

[Applying Crypto Maps to Interfaces](#) 27-20

[Using Interface Access Lists](#) 27-20

[Changing IPsec SA Lifetimes](#) 27-22

[Creating a Basic IPsec Configuration](#) 27-22

[Using Dynamic Crypto Maps](#) 27-24

[Providing Site-to-Site Redundancy](#) 27-26

[Viewing an IPsec Configuration](#) 27-26

[Clearing Security Associations](#) 27-27

[Clearing Crypto Map Configurations](#) 27-27

[Supporting the Nokia VPN Client](#) 27-28

CHAPTER 28

Configuring L2TP over IPSec 28-1

[L2TP Overview](#) 28-1

IPSec Transport and Tunnel Modes	28-2
Configuring L2TP over IPSec Connections	28-2
Tunnel Group Switching	28-5
Viewing L2TP over IPSec Connection Information	28-5
Using L2TP Debug Commands	28-7
Enabling IPSec Debug	28-7
Getting Additional Information	28-8

CHAPTER 29

Setting General IPSec VPN Parameters	29-1
Configuring VPNs in Single, Routed Mode	29-1
Configuring IPSec to Bypass ACLs	29-1
Permitting Intra-Interface Traffic	29-2
NAT Considerations for Intra-Interface Traffic	29-3
Setting Maximum Active IPSec VPN Sessions	29-3
Using Client Update to Ensure Acceptable Client Revision Levels	29-3
Understanding Load Balancing	29-5
Implementing Load Balancing	29-6
Prerequisites	29-6
Eligible Platforms	29-7
Eligible Clients	29-7
VPN Load-Balancing Cluster Configurations	29-7
Some Typical Mixed Cluster Scenarios	29-8
Scenario 1: Mixed Cluster with No WebVPN Connections	29-8
Scenario 2: Mixed Cluster Handling WebVPN Connections	29-8
Configuring Load Balancing	29-9
Configuring the Public and Private Interfaces for Load Balancing	29-9
Configuring the Load Balancing Cluster Attributes	29-10
Configuring VPN Session Limits	29-11

CHAPTER 30

Configuring Tunnel Groups, Group Policies, and Users	30-1
Overview of Tunnel Groups, Group Policies, and Users	30-1
Tunnel Groups	30-2
General Tunnel-Group Connection Parameters	30-2
IPSec Tunnel-Group Connection Parameters	30-3
WebVPN Tunnel-Group Connection Parameters	30-4
Configuring Tunnel Groups	30-5
Maximum Tunnel Groups	30-5
Default IPSec Remote Access Tunnel Group Configuration	30-5

Configuring IPsec Tunnel-Group General Attributes	30-6
Configuring IPsec Remote-Access Tunnel Groups	30-6
Specifying a Name and Type for the IPsec Remote Access Tunnel Group	30-6
Configuring IPsec Remote-Access Tunnel Group General Attributes	30-7
Configuring IPsec Remote-Access Tunnel Group IPsec Attributes	30-10
Configuring IPsec Remote-Access Tunnel Group PPP Attributes	30-12
Configuring LAN-to-LAN Tunnel Groups	30-13
Default LAN-to-LAN Tunnel Group Configuration	30-13
Specifying a Name and Type for a LAN-to-LAN Tunnel Group	30-14
Configuring LAN-to-LAN Tunnel Group General Attributes	30-14
Configuring LAN-to-LAN IPsec Attributes	30-15
Configuring WebVPN Tunnel Groups	30-17
Specifying a Name and Type for a WebVPN Tunnel Group	30-17
Configuring WebVPN Tunnel-Group General Attributes	30-17
Configuring WebVPN Tunnel-Group WebVPN Attributes	30-20
Customizing Login Windows for WebVPN Users	30-23
Configuring Microsoft Active Directory Settings for Password Management	30-24
Using Active Directory to Force the User to Change Password at Next Logon	30-25
Using Active Directory to Specify Maximum Password Age	30-27
Using Active Directory to Override an Account Disabled AAA Indicator	30-28
Using Active Directory to Enforce Minimum Password Length	30-29
Using Active Directory to Enforce Password Complexity	30-30
Group Policies	30-31
Default Group Policy	30-32
Configuring Group Policies	30-34
Configuring an External Group Policy	30-34
Configuring an Internal Group Policy	30-35
Configuring Group Policy Attributes	30-35
Configuring WINS and DNS Servers	30-35
Configuring VPN-Specific Attributes	30-36
Configuring Security Attributes	30-39
Configuring the Banner Message	30-41
Configuring IPsec-UDP Attributes	30-41
Configuring Split-Tunneling Attributes	30-42
Configuring Domain Attributes for Tunneling	30-43
Configuring Attributes for VPN Hardware Clients	30-45
Configuring Backup Server Attributes	30-48
Configuring Microsoft Internet Explorer Client Parameters	30-49
Configuring Network Admission Control Parameters	30-51
Configuring Address Pools	30-54

Configuring Firewall Policies	30-55
Configuring Client Access Rules	30-58
Configuring Group-Policy WebVPN Attributes	30-59
Configuring User Attributes	30-70
Viewing the Username Configuration	30-71
Configuring Attributes for Specific Users	30-71
Setting a User Password and Privilege Level	30-71
Configuring User Attributes	30-72
Configuring VPN User Attributes	30-72
Configuring WebVPN for Specific Users	30-76

CHAPTER 31**Configuring IP Addresses for VPNs 31-1**

Configuring an IP Address Assignment Method	31-1
Configuring Local IP Address Pools	31-2
Configuring AAA Addressing	31-2
Configuring DHCP Addressing	31-3

CHAPTER 32**Configuring Remote Access IPSec VPNs 32-1**

Summary of the Configuration	32-1
Configuring Interfaces	32-2
Configuring ISAKMP Policy and Enabling ISAKMP on the Outside Interface	32-3
Configuring an Address Pool	32-4
Adding a User	32-4
Creating a Transform Set	32-4
Defining a Tunnel Group	32-5
Creating a Dynamic Crypto Map	32-6
Creating a Crypto Map Entry to Use the Dynamic Crypto Map	32-7

CHAPTER 33**Configuring Network Admission Control 33-1**

Uses, Requirements, and Limitations	33-1
Configuring Basic Settings	33-1
Specifying the Access Control Server Group	33-2
Enabling NAC	33-2
Configuring the Default ACL for NAC	33-3
Configuring Exemptions from NAC	33-4
Changing Advanced Settings	33-5
Changing Clientless Authentication Settings	33-5
Enabling and Disabling Clientless Authentication	33-5

Changing the Login Credentials Used for Clientless Authentication	33-6
Configuring NAC Session Attributes	33-7
Setting the Query-for-Posture-Changes Timer	33-8
Setting the Revalidation Timer	33-9

CHAPTER 34

Configuring Easy VPN Services on the ASA 5505	34-1
Specifying the Client/Server Role of the Cisco ASA 5505	34-1
Specifying the Primary and Secondary Servers	34-2
Specifying the Mode	34-3
NEM with Multiple Interfaces	34-3
Configuring Automatic Xauth Authentication	34-4
Configuring IPSec Over TCP	34-4
Comparing Tunneling Options	34-5
Specifying the Tunnel Group or Trustpoint	34-6
Specifying the Tunnel Group	34-6
Specifying the Trustpoint	34-7
Configuring Split Tunneling	34-7
Configuring Device Pass-Through	34-8
Configuring Remote Management	34-8
Guidelines for Configuring the Easy VPN Server	34-9
Group Policy and User Attributes Pushed to the Client	34-9
Authentication Options	34-11

CHAPTER 35

Configuring the PPPoE Client	35-1
PPPoE Client Overview	35-1
Configuring the PPPoE Client Username and Password	35-2
Enabling PPPoE	35-3
Using PPPoE with a Fixed IP Address	35-3
Monitoring and Debugging the PPPoE Client	35-4
Clearing the Configuration	35-5
Using Related Commands	35-5

CHAPTER 36

Configuring LAN-to-LAN IPsec VPNs	36-1
Summary of the Configuration	36-1
Configuring Interfaces	36-2
Configuring ISAKMP Policy and Enabling ISAKMP on the Outside Interface	36-2
Creating a Transform Set	36-4

Configuring an ACL	36-4
Defining a Tunnel Group	36-5
Creating a Crypto Map and Applying It To an Interface	36-6
Applying Crypto Maps to Interfaces	36-7

CHAPTER 37**Configuring WebVPN 37-1**

Getting Started with WebVPN	37-1
Observing WebVPN Security Precautions	37-2
Understanding Features Not Supported for WebVPN	37-2
Using SSL to Access the Central Site	37-3
Using HTTPS for WebVPN Sessions	37-3
Configuring WebVPN and ASDM on the Same Interface	37-3
Setting WebVPN HTTP/HTTPS Proxy	37-4
Configuring SSL/TLS Encryption Protocols	37-4
Authenticating with Digital Certificates	37-5
Enabling Cookies on Browsers for WebVPN	37-5
Managing Passwords	37-5
Using Single Sign-on with WebVPN	37-6
Configuring SSO with HTTP Basic or NTLM Authentication	37-6
Configuring SSO Authentication Using SiteMinder	37-7
Configuring SSO with the HTTP Form Protocol	37-9
Authenticating with Digital Certificates	37-15
Creating and Applying WebVPN Policies	37-15
Creating Port Forwarding, URL, and Access Lists in Global Configuration Mode	37-16
Assigning Lists to Group Policies and Users in Group-Policy or User Mode	37-16
Enabling Features for Group Policies and Users	37-16
Assigning Users to Group Policies	37-16
Using the Security Appliance Authentication Server	37-16
Using a RADIUS Server	37-16
Configuring WebVPN Tunnel Group Attributes	37-17
Configuring WebVPN Group Policy and User Attributes	37-17
Configuring Application Access	37-18
Downloading the Port-Forwarding Applet Automatically	37-18
Closing Application Access to Prevent hosts File Errors	37-18
Recovering from hosts File Errors When Using Application Access	37-18
Understanding the hosts File	37-19
Stopping Application Access Improperly	37-19
Reconfiguring a hosts File	37-20
Configuring File Access	37-22

Configuring Access to Citrix MetaFrame Services	37-24
Using WebVPN with PDAs	37-25
Using E-Mail over WebVPN	37-26
Configuring E-mail Proxies	37-26
E-mail Proxy Certificate Authentication	37-27
Configuring MAPI	37-27
Configuring Web E-mail: MS Outlook Web Access	37-27
Optimizing WebVPN Performance	37-28
Configuring Caching	37-28
Configuring Content Transformation	37-28
Configuring a Certificate for Signing Rewritten Java Content	37-29
Disabling Content Rewrite	37-29
Using Proxy Bypass	37-29
Configuring Application Profile Customization Framework	37-30
APCF Syntax	37-30
APCF Example	37-32
WebVPN End User Setup	37-32
Defining the End User Interface	37-32
Viewing the WebVPN Home Page	37-33
Viewing the WebVPN Application Access Panel	37-33
Viewing the Floating Toolbar	37-34
Customizing WebVPN Pages	37-35
Using Cascading Style Sheet Parameters	37-35
Customizing the WebVPN Login Page	37-36
Customizing the WebVPN Logout Page	37-37
Customizing the WebVPN Home Page	37-38
Customizing the Application Access Window	37-40
Customizing the Prompt Dialogs	37-41
Applying Customizations to Tunnel Groups, Groups and Users	37-42
Requiring Usernames and Passwords	37-43
Communicating Security Tips	37-44
Configuring Remote Systems to Use WebVPN Features	37-44
Capturing WebVPN Data	37-50
Creating a Capture File	37-51
Using a Browser to Display Capture Data	37-51

CHAPTER 38

Configuring SSL VPN Client 38-1

Installing SVC	38-1
Platform Requirements	38-1

Installing the SVC Software	38-2
Enabling SVC	38-3
Enabling Permanent SVC Installation	38-4
Enabling Rekey	38-5
Enabling and Adjusting Dead Peer Detection	38-5
Enabling Keepalive	38-6
Using SVC Compression	38-6
Viewing SVC Sessions	38-7
Logging Off SVC Sessions	38-8
Updating SVCs	38-8

CHAPTER 39

Configuring Certificates	39-1
Public Key Cryptography	39-1
About Public Key Cryptography	39-1
Certificate Scalability	39-2
About Key Pairs	39-2
About Trustpoints	39-3
About Revocation Checking	39-3
About CRLs	39-3
About OCSP	39-4
Supported CA Servers	39-5
Certificate Configuration	39-5
Preparing for Certificates	39-5
Configuring Key Pairs	39-6
Generating Key Pairs	39-6
Removing Key Pairs	39-7
Configuring Trustpoints	39-7
Obtaining Certificates	39-9
Obtaining Certificates with SCEP	39-9
Obtaining Certificates Manually	39-11
Configuring CRLs for a Trustpoint	39-13
Exporting and Importing Trustpoints	39-14
Exporting a Trustpoint Configuration	39-15
Importing a Trustpoint Configuration	39-15
Configuring CA Certificate Map Rules	39-15

PART 4

System Administration

CHAPTER 40**Managing System Access 40-1**

Allowing Telnet Access 40-1

Allowing SSH Access 40-2

Configuring SSH Access 40-2

Using an SSH Client 40-3

Allowing HTTPS Access for ASDM 40-3

Configuring ASDM and WebVPN on the Same Interface 40-4

Configuring AAA for System Administrators 40-5

Configuring Authentication for CLI Access 40-5

Configuring Authentication To Access Privileged EXEC Mode 40-6

Configuring Authentication for the Enable Command 40-6

Authenticating Users Using the Login Command 40-6

Configuring Command Authorization 40-7

Command Authorization Overview 40-7

Configuring Local Command Authorization 40-8

Configuring TACACS+ Command Authorization 40-11

Configuring Command Accounting 40-14

Viewing the Current Logged-In User 40-14

Recovering from a Lockout 40-15

Configuring a Login Banner 40-16

CHAPTER 41**Managing Software, Licenses, and Configurations 41-1**

Managing Licenses 41-1

Obtaining an Activation Key 41-1

Entering a New Activation Key 41-2

Viewing Files in Flash Memory 41-2

Retrieving Files from Flash Memory 41-3

Downloading Software or Configuration Files to Flash Memory 41-3

Downloading a File to a Specific Location 41-4

Downloading a File to the Startup or Running Configuration 41-4

Configuring the Application Image and ASDM Image to Boot 41-5

Configuring the File to Boot as the Startup Configuration 41-6

Performing Zero Downtime Upgrades for Failover Pairs 41-6

Upgrading an Active/Standby Failover Configuration 41-7

Upgrading and Active/Active Failover Configuration 41-8

Backing Up Configuration Files 41-8

Backing up the Single Mode Configuration or Multiple Mode System Configuration 41-9

Backing Up a Context Configuration in Flash Memory 41-9

Backing Up a Context Configuration within a Context	41-9
Copying the Configuration from the Terminal Display	41-10
Configuring Auto Update Support	41-10
Configuring Communication with an Auto Update Server	41-10
Configuring Client Updates as an Auto Update Server	41-12
Viewing Auto Update Status	41-13

CHAPTER 42**Monitoring the Security Appliance 42-1**

Using SNMP	42-1
SNMP Overview	42-1
Enabling SNMP	42-3
Configuring and Managing Logs	42-5
Logging Overview	42-5
Logging in Multiple Context Mode	42-5
Enabling and Disabling Logging	42-6
Enabling Logging to All Configured Output Destinations	42-6
Disabling Logging to All Configured Output Destinations	42-6
Viewing the Log Configuration	42-6
Configuring Log Output Destinations	42-7
Sending System Log Messages to a Syslog Server	42-7
Sending System Log Messages to the Console Port	42-8
Sending System Log Messages to an E-mail Address	42-9
Sending System Log Messages to ASDM	42-10
Sending System Log Messages to a Telnet or SSH Session	42-11
Sending System Log Messages to the Log Buffer	42-12
Filtering System Log Messages	42-14
Message Filtering Overview	42-15
Filtering System Log Messages by Class	42-15
Filtering System Log Messages with Custom Message Lists	42-17
Customizing the Log Configuration	42-18
Customizing the Log Configuration	42-18
Configuring the Logging Queue	42-19
Including the Date and Time in System Log Messages	42-19
Including the Device ID in System Log Messages	42-19
Generating System Log Messages in EMBLEM Format	42-20
Disabling a System Log Message	42-20
Changing the Severity Level of a System Log Message	42-21
Changing the Amount of Internal Flash Memory Available for Logs	42-22
Understanding System Log Messages	42-23

System Log Message Format	42-23
Severity Levels	42-23

CHAPTER 43

Troubleshooting the Security Appliance 43-1

Testing Your Configuration	43-1
Enabling ICMP Debug Messages and System Messages	43-1
Pinging Security Appliance Interfaces	43-2
Pinging Through the Security Appliance	43-4
Disabling the Test Configuration	43-5
Traceroute	43-6
Packet Tracer	43-6
Reloading the Security Appliance	43-6
Performing Password Recovery	43-7
Performing Password Recovery for the ASA 5500 Series Adaptive Security Appliance	43-7
Password Recovery for the PIX 500 Series Security Appliance	43-8
Disabling Password Recovery	43-9
Resetting the Password on the SSM Hardware Module	43-10
Other Troubleshooting Tools	43-10
Viewing Debug Messages	43-11
Capturing Packets	43-11
Viewing the Crash Dump	43-11
Common Problems	43-11

PART 2

Reference

Supported Platforms and Feature Licenses	A-1
Security Services Module Support	A-9
VPN Specifications	A-10
Cisco VPN Client Support	A-11
Cisco Secure Desktop Support	A-11
Site-to-Site VPN Compatibility	A-11
Cryptographic Standards	A-12
Example 1: Multiple Mode Firewall With Outside Access	B-1
Example 1: System Configuration	B-2
Example 1: Admin Context Configuration	B-4
Example 1: Customer A Context Configuration	B-4
Example 1: Customer B Context Configuration	B-4
Example 1: Customer C Context Configuration	B-5
Example 2: Single Mode Firewall Using Same Security Level	B-6

Example 3: Shared Resources for Multiple Contexts	B-8
Example 3: System Configuration	B-9
Example 3: Admin Context Configuration	B-9
Example 3: Department 1 Context Configuration	B-10
Example 3: Department 2 Context Configuration	B-11
Example 4: Multiple Mode, Transparent Firewall with Outside Access	B-12
Example 4: System Configuration	B-13
Example 4: Admin Context Configuration	B-14
Example 4: Customer A Context Configuration	B-15
Example 4: Customer B Context Configuration	B-15
Example 4: Customer C Context Configuration	B-16
Example 5: WebVPN Configuration	B-16
Example 6: IPv6 Configuration	B-18
Example 7: Cable-Based Active/Standby Failover (Routed Mode)	B-20
Example 8: LAN-Based Active/Standby Failover (Routed Mode)	B-21
Example 8: Primary Unit Configuration	B-21
Example 8: Secondary Unit Configuration	B-22
Example 9: LAN-Based Active/Active Failover (Routed Mode)	B-22
Example 9: Primary Unit Configuration	B-23
Example 9: Primary System Configuration	B-23
Example 9: Primary admin Context Configuration	B-24
Example 9: Primary ctx1 Context Configuration	B-25
Example 9: Secondary Unit Configuration	B-25
Example 10: Cable-Based Active/Standby Failover (Transparent Mode)	B-26
Example 11: LAN-Based Active/Standby Failover (Transparent Mode)	B-27
Example 11: Primary Unit Configuration	B-27
Example 11: Secondary Unit Configuration	B-28
Example 12: LAN-Based Active/Active Failover (Transparent Mode)	B-28
Example 12: Primary Unit Configuration	B-29
Example 12: Primary System Configuration	B-29
Example 12: Primary admin Context Configuration	B-30
Example 12: Primary ctx1 Context Configuration	B-31
Example 12: Secondary Unit Configuration	B-31
Example 13: Dual ISP Support Using Static Route Tracking	B-31
Example 14: ASA 5505 Base License	B-33
Example 15: ASA 5505 Security Plus License with Failover and Dual-ISP Backup	B-35
Example 15: Primary Unit Configuration	B-35
Example 15: Secondary Unit Configuration	B-37

Example 16: Network Traffic Diversion	B-37
Inspecting All Traffic with the AIP SSM	B-43
Inspecting Specific Traffic with the AIP SSM	B-44
Verifying the Recording of Alert Events	B-45
Troubleshooting the Configuration	B-47
Firewall Mode and Security Context Mode	C-1
Command Modes and Prompts	C-2
Syntax Formatting	C-3
Abbreviating Commands	C-3
Command-Line Editing	C-3
Command Completion	C-4
Command Help	C-4
Filtering show Command Output	C-4
Command Output Paging	C-5
Adding Comments	C-6
Text Configuration Files	C-6
How Commands Correspond with Lines in the Text File	C-6
Command-Specific Configuration Mode Commands	C-6
Automatic Text Entries	C-7
Line Order	C-7
Commands Not Included in the Text Configuration	C-7
Passwords	C-7
Multiple Security Context Files	C-7
IPv4 Addresses and Subnet Masks	D-1
Classes	D-1
Private Networks	D-2
Subnet Masks	D-2
Determining the Subnet Mask	D-3
Determining the Address to Use with the Subnet Mask	D-3
IPv6 Addresses	D-5
IPv6 Address Format	D-5
IPv6 Address Types	D-6
Unicast Addresses	D-6
Multicast Address	D-8
Anycast Address	D-9
Required Addresses	D-10
IPv6 Address Prefixes	D-10
Protocols and Applications	D-11

TCP and UDP Ports	D-11
Local Ports and Protocols	D-14
ICMP Types	D-15
Selecting LDAP, RADIUS, or Local Authentication and Authorization	E-1
Understanding Policy Enforcement of Permissions and Attributes	E-2
Configuring an External LDAP Server	E-2
Reviewing the LDAP Directory Structure and Configuration Procedure	E-3
Organizing the Security Appliance LDAP Schema	E-3
Searching the Hierarchy	E-4
Binding the Security Appliance to the LDAP Server	E-5
Defining the Security Appliance LDAP Schema	E-5
Cisco -AV-Pair Attribute Syntax	E-14
Example Security Appliance Authorization Schema	E-15
Loading the Schema in the LDAP Server	E-18
Defining User Permissions	E-18
Example User File	E-18
Reviewing Examples of Active Directory Configurations	E-19
Example 1: Configuring LDAP Authorization with Microsoft Active Directory (ASA/PIX)	E-19
Example 2: Configuring LDAP Authentication with Microsoft Active Directory	E-20
Example 3: LDAP Authentication and LDAP Authorization with Microsoft Active Directory	E-22
Configuring an External RADIUS Server	E-24
Reviewing the RADIUS Configuration Procedure	E-24
Security Appliance RADIUS Authorization Attributes	E-25
Security Appliance TACACS+ Attributes	E-32

GLOSSARY

INDEX



About This Guide

This preface introduces the *Cisco Security Appliance Command Line Configuration Guide*, and includes the following sections:

- [Document Objectives, page xxxv](#)
- [Audience, page xxxv](#)
- [Related Documentation, page xxxvi](#)
- [Document Organization, page xxxvi](#)
- [Document Conventions, page xxxix](#)
- [, page xxxix](#)

Document Objectives

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

You can also configure and monitor the security appliance by using ASDM, a web-based GUI application. ASDM includes configuration wizards to guide you through some common configuration scenarios, and online Help for less common scenarios. For more information, see: <http://www.cisco.com/univercd/cc/td/doc/product/netsec/secmgmt/asdm/index.htm>

This guide applies to the Cisco PIX 500 series security appliances (PIX 515E, PIX 525, and PIX 535) and the Cisco ASA 5500 series security appliances (ASA 5505, ASA 5510, ASA 5520, ASA 5540, and ASA 5550). Throughout this guide, the term “security appliance” applies generically to all supported models, unless specified otherwise. The PIX 501, PIX 506E, and PIX 520 security appliances are not supported.

Audience

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

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

Related Documentation

For more information, refer to the following documentation:

- *Cisco PIX Security Appliance Release Notes*
- *Cisco ASDM Release Notes*
- *Cisco PIX 515E Quick Start Guide*
- *Guide for Cisco PIX 6.2 and 6.3 Users Upgrading to Cisco PIX Software Version 7.0*
- *Migrating to ASA for VPN 3000 Series Concentrator Administrators*
- *Cisco Security Appliance Command Reference*
- *Cisco ASA 5500 Series Adaptive Security Appliance Getting Started Guide*
- *Cisco ASA 5500 Series Release Notes*
- *Cisco Security Appliance Logging Configuration and System Log Messages*
- *Cisco Secure Desktop Configuration Guide for Cisco ASA 5500 Series Administrators*

Document Organization

This guide includes the chapters and appendixes described in [Table 1](#).

Table 1 **Document Organization**

Chapter/Appendix	Definition
Part 1: Getting Started and General Information	
Chapter 1, “Introduction to the Security Appliance”	Provides a high-level overview of the security appliance.
Chapter 2, “Getting Started”	Describes how to access the command-line interface, configure the firewall mode, and work with the configuration.
Chapter 3, “Enabling Multiple Context Mode”	Describes how to use security contexts and enable multiple context mode.
Chapter 4, “Configuring Switch Ports and VLAN Interfaces for the Cisco ASA 5505 Adaptive Security Appliance”	Describes how to configure switch ports and VLAN interfaces for the ASA 5505 adaptive security appliance.
Chapter 5, “Configuring Ethernet Settings and Subinterfaces”	Describes how to configure Ethernet settings for physical interfaces and add subinterfaces.
Chapter 6, “Adding and Managing Security Contexts”	Describes how to configure multiple security contexts on the security appliance.
Chapter 7, “Configuring Interface Parameters”	Describes how to configure each interface and subinterface for a name, security, level, and IP address.
Chapter 8, “Configuring Basic Settings”	Describes how to configure basic settings that are typically required for a functioning configuration.
Chapter 9, “Configuring IP Routing”	Describes how to configure IP routing.

Table 1 **Document Organization (continued)**

Chapter/Appendix	Definition
Chapter 10, “Configuring DHCP, DDNS, and WCCP Services”	Describes how to configure the DHCP server and DHCP relay.
Chapter 11, “Configuring Multicast Routing”	Describes how to configure multicast routing.
Chapter 12, “Configuring IPv6”	Describes how to enable and configure IPv6.
Chapter 13, “Configuring AAA Servers and the Local Database”	Describes how to configure AAA servers and the local database.
Chapter 14, “Configuring Failover”	Describes the failover feature, which lets you configure two security appliances so that one will take over operation if the other one fails.
Part 2: Configuring the Firewall	
Chapter 15, “Firewall Mode Overview”	Describes in detail the two operation modes of the security appliance, routed and transparent mode, and how data is handled differently with each mode.
Chapter 16, “Identifying Traffic with Access Lists”	Describes how to identify traffic with access lists.
Chapter 17, “Applying NAT”	Describes how address translation is performed.
Chapter 18, “Permitting or Denying Network Access”	Describes how to control network access through the security appliance using access lists.
Chapter 19, “Applying AAA for Network Access”	Describes how to enable AAA for network access.
Chapter 20, “Applying Filtering Services”	Describes ways to filter web traffic to reduce security risks or prevent inappropriate use.
Chapter 21, “Using Modular Policy Framework”	Describes how to use the Modular Policy Framework to create security policies for TCP, general connection settings, inspection, and QoS.
Chapter 22, “Managing AIP SSM and CSC SSM”	Describes how to configure the security appliance to send traffic to an AIP SSM or a CSC SSM, how to check the status of an SSM, and how to update the software image on an intelligent SSM.
Chapter 23, “Preventing Network Attacks”	Describes how to configure protection features to intercept and respond to network attacks.
Chapter 24, “Configuring QoS”	Describes how to configure the network to provide better service to selected network traffic over various technologies, including Frame Relay, Asynchronous Transfer Mode (ATM), Ethernet and 802.1 networks, SONET, and IP routed networks.
Chapter 25, “Configuring Application Layer Protocol Inspection”	Describes how to use and configure application inspection.
Chapter 26, “Configuring ARP Inspection and Bridging Parameters”	Describes how to enable ARP inspection and how to customize bridging operations.
Part 3: Configuring VPN	
Chapter 27, “Configuring IPsec and ISAKMP”	Describes how to configure ISAKMP and IPsec tunneling to build and manage VPN “tunnels,” or secure connections between remote users and a private corporate network.

Table 1 **Document Organization (continued)**

Chapter/Appendix	Definition
Chapter 28, “Configuring L2TP over IPsec”	Describes how to configure IPsec over L2TP on the security appliance.
Chapter 29, “Setting General IPsec VPN Parameters”	Describes miscellaneous VPN configuration procedures.
Chapter 30, “Configuring Tunnel Groups, Group Policies, and Users”	Describes how to configure VPN tunnel groups, group policies, and users.
Chapter 31, “Configuring IP Addresses for VPNs”	Describes how to configure IP addresses in your private network addressing scheme, which let the client function as a tunnel endpoint.
Chapter 32, “Configuring Remote Access IPsec VPNs”	Describes how to configure a remote access VPN connection.
Chapter 33, “Configuring Network Admission Control”	Describes how to configure Network Admission Control (NAC).
Chapter 34, “Configuring Easy VPN Services on the ASA 5505”	Describes how to configure Easy VPN on the ASA 5505 adaptive security appliance.
Chapter 35, “Configuring the PPPoE Client”	Describes how to configure the PPPoE client provided with the security appliance.
Chapter 36, “Configuring LAN-to-LAN IPsec VPNs”	Describes how to build a LAN-to-LAN VPN connection.
Chapter 37, “Configuring WebVPN”	Describes how to establish a secure, remote-access VPN tunnel to a security appliance using a web browser.
Chapter 38, “Configuring SSL VPN Client”	Describes how to install and configure the SSL VPN Client.
Chapter 39, “Configuring Certificates”	Describes how to configure a digital certificates, which contains information that identifies a user or device. Such information can include a name, serial number, company, department, or IP address. A digital certificate also contains a copy of the public key for the user or device.
Part 4: System Administration	
Chapter 40, “Managing System Access”	Describes how to access the security appliance for system management through Telnet, SSH, and HTTPS.
Chapter 41, “Managing Software, Licenses, and Configurations”	Describes how to enter license keys and download software and configurations files.
Chapter 42, “Monitoring the Security Appliance”	Describes how to monitor the security appliance.
Chapter 43, “Troubleshooting the Security Appliance”	Describes how to troubleshoot the security appliance.
Part 4: Reference	
Appendix A, “Feature Licenses and Specifications”	Describes the feature licenses and specifications.
Appendix B, “Sample Configurations”	Describes a number of common ways to implement the security appliance.

Table 1 **Document Organization (continued)**

Chapter/Appendix	Definition
Appendix C, “Using the Command-Line Interface”	Describes how to use the CLI to configure the the security appliance.
Appendix D, “Addresses, Protocols, and Ports”	Provides a quick reference for IP addresses, protocols, and applications.
Appendix E, “Configuring an External Server for Authorization and Authentication”	Provides information about configuring LDAP and RADIUS authorization servers.
“Glossary”	Provides a handy reference for commonly-used terms and acronyms.
“Index”	Provides an index for the guide.

Document Conventions

Command descriptions use these conventions:

- Braces ({ }) indicate a required choice.
- Square brackets ([]) indicate optional elements.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- **Boldface** indicates commands and keywords that are entered literally as shown.
- *Italics* indicate arguments for which you supply values.

Examples use these conventions:

- Examples depict screen displays and the command line in *screen* font.
- Information you need to enter in examples is shown in **boldface screen** font.
- Variables for which you must supply a value are shown in *italic screen* font.



Note

Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Obtaining Documentation and Submitting a Service Request

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

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

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



PART 1

Getting Started and General Information



CHAPTER 1

Introduction to the Security Appliance

The security appliance combines advanced stateful firewall and VPN concentrator functionality in one device, and for some models, an integrated intrusion prevention module called the AIP SSM or an integrated content security and control module called the CSC SSM. The security appliance includes many advanced features, such as multiple security contexts (similar to virtualized firewalls), transparent (Layer 2) firewall or routed (Layer 3) firewall operation, advanced inspection engines, IPSec and WebVPN support, and many more features. See [Appendix A, “Feature Licenses and Specifications,”](#) for a list of supported platforms and features. For a list of new features, see the *Cisco ASA 5500 Series Release Notes* or the *Cisco PIX Security Appliance Release Notes*.



Note

The Cisco PIX 501 and PIX 506E security appliances are not supported.

This chapter includes the following sections:

- [Firewall Functional Overview, page 1-1](#)
- [VPN Functional Overview, page 1-5](#)
- [Intrusion Prevention Services Functional Overview, page 1-5](#)
- [Security Context Overview, page 1-6](#)

Firewall Functional Overview

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

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

This section includes the following topics:

- [Security Policy Overview, page 1-2](#)
- [Firewall Mode Overview, page 1-3](#)
- [Stateful Inspection Overview, page 1-4](#)

Security Policy Overview

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

- [Permitting or Denying Traffic with Access Lists, page 1-2](#)
- [Applying NAT, page 1-2](#)
- [Using AAA for Through Traffic, page 1-2](#)
- [Applying HTTP, HTTPS, or FTP Filtering, page 1-3](#)
- [Applying Application Inspection, page 1-3](#)
- [Sending Traffic to the Advanced Inspection and Prevention Security Services Module, page 1-3](#)
- [Sending Traffic to the Content Security and Control Security Services Module, page 1-3](#)
- [Applying QoS Policies, page 1-3](#)
- [Applying Connection Limits and TCP Normalization, page 1-3](#)

Permitting or Denying Traffic with Access Lists

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

Applying NAT

Some of the benefits of NAT include the following:

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

Using AAA for Through Traffic

You can require authentication and/or authorization for certain types of traffic, for example, for HTTP. The security appliance also sends accounting information to a RADIUS or TACACS+ server.

Applying HTTP, HTTPS, or FTP Filtering

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

- Websense Enterprise
- Secure Computing SmartFilter

Applying Application Inspection

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

Sending Traffic to the Advanced Inspection and Prevention Security Services Module

If your model supports the AIP SSM for intrusion prevention, then you can send traffic to the AIP SSM for inspection.

Sending Traffic to the Content Security and Control Security Services Module

If your model supports it, the CSC SSM provides protection against viruses, spyware, spam, and other unwanted traffic. It accomplishes this by scanning the FTP, HTTP, POP3, and SMTP traffic that you configure the adaptive security appliance to send to it.

Applying QoS Policies

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

Applying Connection Limits and TCP Normalization

You can limit TCP and UDP connections and embryonic connections. Limiting the number of connections and embryonic connections protects you from a DoS attack. The security appliance uses the embryonic limit to trigger TCP Intercept, which protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination.

TCP normalization is a feature consisting of advanced TCP connection settings designed to drop packets that do not appear normal.

Firewall Mode Overview

The security appliance runs in two different firewall modes:

- Routed
- Transparent

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

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

You might use a transparent firewall to simplify your network configuration. Transparent mode is also useful if you want the firewall to be invisible to attackers. You can also use a transparent firewall for traffic that would otherwise be blocked in routed mode. For example, a transparent firewall can allow multicast streams using an EtherType access list.

Stateful Inspection Overview

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

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

- Is this a new connection?

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

The session management path is responsible for the following tasks:

- Performing the access list checks
- Performing route lookups
- Allocating NAT translations (xlates)
- Establishing sessions in the “fast path”



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

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

- Is this an established connection?

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

- IP checksum verification
- Session lookup
- TCP sequence number check
- NAT translations based on existing sessions
- Layer 3 and Layer 4 header adjustments

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

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

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

VPN Functional Overview

A VPN is a secure connection across a TCP/IP network (such as the Internet) that appears as a private connection. This secure connection is called a tunnel. The security appliance uses tunneling protocols to negotiate security parameters, create and manage tunnels, encapsulate packets, transmit or receive them through the tunnel, and unencapsulate them. The security appliance functions as a bidirectional tunnel endpoint: it can receive plain packets, encapsulate them, and send them to the other end of the tunnel where they are unencapsulated and sent to their final destination. It can also receive encapsulated packets, unencapsulate them, and send them to their final destination. The security appliance invokes various standard protocols to accomplish these functions.

The security appliance performs the following functions:

- Establishes tunnels
- Negotiates tunnel parameters
- Authenticates users
- Assigns user addresses
- Encrypts and decrypts data
- Manages security keys
- Manages data transfer across the tunnel
- Manages data transfer inbound and outbound as a tunnel endpoint or router

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

Intrusion Prevention Services Functional Overview

The Cisco ASA 5500 series adaptive security appliance supports the AIP SSM, an intrusion prevention services module that monitors and performs real-time analysis of network traffic by looking for anomalies and misuse based on an extensive, embedded signature library. When the system detects unauthorized activity, it can terminate the specific connection, permanently block the attacking host, log the incident, and send an alert to the device manager. Other legitimate connections continue to operate independently without interruption. For more information, see *Configuring the Cisco Intrusion Prevention System Sensor Using the Command Line Interface*.

Security Context Overview

You can partition a single security appliance into multiple virtual devices, known as security contexts. Each context is an independent device, with its own security policy, interfaces, and administrators. Multiple contexts are similar to having multiple standalone devices. Many features are supported in multiple context mode, including routing tables, firewall features, IPS, and management. Some features are not supported, including VPN and dynamic routing protocols.

In multiple context mode, the security appliance includes a configuration for each context that identifies the security policy, interfaces, and almost all the options you can configure on a standalone device. The system administrator adds and manages contexts by configuring them in the system configuration, which, like a single mode configuration, is the startup configuration. The system configuration identifies basic settings for the security appliance. The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses one of the contexts that is designated as the admin context.

The admin context is just like any other context, except that when a user logs into the admin context, then that user has system administrator rights and can access the system and all other contexts.

**Note**

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

Multiple context mode supports static routing only.



CHAPTER 2

Getting Started

This chapter describes how to access the command-line interface, configure the firewall mode, and work with the configuration. This chapter includes the following sections:

- [Getting Started with Your Platform Model, page 2-1](#)
- [Factory Default Configurations, page 2-1](#)
- [Accessing the Command-Line Interface, page 2-4](#)
- [Setting Transparent or Routed Firewall Mode, page 2-5](#)
- [Working with the Configuration, page 2-6](#)

Getting Started with Your Platform Model

This guide applies to multiple security appliance platforms and models: the PIX 500 series security appliances and the ASA 5500 series adaptive security appliances. There are some hardware differences between the PIX and the ASA security appliance. Moreover, the ASA 5505 includes a built-in switch, and requires some special configuration. For these hardware-based differences, the platforms or models supported are noted directly in each section.

Some models do not support all features covered in this guide. For example, the ASA 5505 adaptive security appliance does not support security contexts. This guide might not list each supported model when discussing a feature. To determine the features that are supported for your model before you start your configuration, see the [“Supported Platforms and Feature Licenses” section on page A-1](#) for a detailed list of the features supported for each model.

Factory Default Configurations

The factory default configuration is the configuration applied by Cisco to new security appliances. The factory default configuration is supported on all models except for the PIX 525 and PIX 535 security appliances.

For the PIX 515/515E and the ASA 5510 and higher security appliances, the factory default configuration configures an interface for management so you can connect to it using ASDM, with which you can then complete your configuration.

For the ASA 5505 adaptive security appliance, the factory default configuration configures interfaces and NAT so that the security appliance is ready to use in your network immediately.

The factory default configuration is available only for routed firewall mode and single context mode. See [Chapter 3, “Enabling Multiple Context Mode,”](#) for more information about multiple context mode. See the [“Setting Transparent or Routed Firewall Mode”](#) section on [page 2-5](#) for more information about routed and transparent firewall mode.

This section includes the following topics:

- [Restoring the Factory Default Configuration, page 2-2](#)
- [ASA 5505 Default Configuration, page 2-2](#)
- [ASA 5510 and Higher Default Configuration, page 2-3](#)
- [PIX 515/515E Default Configuration, page 2-4](#)

Restoring the Factory Default Configuration

To restore the factory default configuration, enter the following command:

```
hostname(config)# configure factory-default [ip_address [mask]]
```

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.

After you restore the factory default configuration, save it to internal Flash memory using the **write memory** command. The **write memory** 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.



Note

This command also clears the **boot system** command, if present, along with the rest of the configuration. The **boot system** command lets you boot from a specific image, including an image on the external Flash memory card. The next time you reload the security appliance 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 security appliance does not boot.

To configure additional settings that are useful for a full configuration, see the **setup** command.

ASA 5505 Default Configuration

The default factory configuration for the ASA 5505 adaptive security appliance configures the following:

- An inside VLAN 1 interface that includes the Ethernet 0/1 through 0/7 switch ports. If you did not set the IP address in the **configure factory-default** command, then the VLAN 1 IP address and mask are 192.168.1.1 and 255.255.255.0.
- An outside VLAN 2 interface that includes the Ethernet 0/0 switch port. VLAN 2 derives its IP address using DHCP.
- The default route is also derived from DHCP.
- All inside IP addresses are translated when accessing the outside using interface PAT.
- By default, inside users can access the outside with an access list, and outside users are prevented from accessing the inside.

- The DHCP server is enabled on the security appliance, so a PC connecting to the VLAN 1 interface receives an address between 192.168.1.2 and 192.168.1.254.
- The HTTP server is enabled for ASDM and is accessible to users on the 192.168.1.0 network.

The configuration consists of the following commands:

```
interface Ethernet 0/0
  switchport access vlan 2
  no shutdown
interface Ethernet 0/1
  switchport access vlan 1
  no shutdown
interface Ethernet 0/2
  switchport access vlan 1
  no shutdown
interface Ethernet 0/3
  switchport access vlan 1
  no shutdown
interface Ethernet 0/4
  switchport access vlan 1
  no shutdown
interface Ethernet 0/5
  switchport access vlan 1
  no shutdown
interface Ethernet 0/6
  switchport access vlan 1
  no shutdown
interface Ethernet 0/7
  switchport access vlan 1
  no shutdown
interface vlan2
  nameif outside
  no shutdown
  ip address dhcp setroute
interface vlan1
  nameif inside
  ip address 192.168.1.1 255.255.255.0
  security-level 100
  no shutdown
global (outside) 1 interface
nat (inside) 1 0 0
http server enable
http 192.168.1.0 255.255.255.0 inside
dhcpd address 192.168.1.2-192.168.1.254 inside
dhcpd auto_config outside
dhcpd enable inside
logging asdm informational
```

ASA 5510 and Higher Default Configuration

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

- The management interface, Management 0/0. If you did not set the IP address in the **configure factory-default** command, then the IP address and mask are 192.168.1.1 and 255.255.255.0.
- The DHCP server is enabled on the security appliance, so a PC connecting to the interface receives an address between 192.168.1.2 and 192.168.1.254.
- The HTTP server is enabled for ASDM and is accessible to users on the 192.168.1.0 network.

The configuration consists of the following commands:

```
interface management 0/0
  ip address 192.168.1.1 255.255.255.0
  nameif management
  security-level 100
  no shutdown
asdm logging informational 100
asdm history enable
http server enable
http 192.168.1.0 255.255.255.0 management
dhcpd address 192.168.1.2-192.168.1.254 management
dhcpd lease 3600
dhcpd ping_timeout 750
dhcpd enable management
```

PIX 515/515E Default Configuration

The default factory configuration for the PIX 515/515E security appliance configures the following:

- The inside Ethernet1 interface. If you did not set the IP address in the **configure factory-default** command, then the IP address and mask are 192.168.1.1 and 255.255.255.0.
- The DHCP server is enabled on the security appliance, so a PC connecting to the interface receives an address between 192.168.1.2 and 192.168.1.254.
- The HTTP server is enabled for ASDM and is accessible to users on the 192.168.1.0 network.

The configuration consists of the following commands:

```
interface ethernet 1
  ip address 192.168.1.1 255.255.255.0
  nameif management
  security-level 100
  no shutdown
asdm logging informational 100
asdm history enable
http server enable
http 192.168.1.0 255.255.255.0 management
dhcpd address 192.168.1.2-192.168.1.254 management
dhcpd lease 3600
dhcpd ping_timeout 750
dhcpd enable management
```

Accessing the Command-Line Interface

For initial configuration, access the command-line interface directly from the console port. Later, you can configure remote access using Telnet or SSH according to [Chapter 40, “Managing System Access.”](#) If your system is already in multiple context mode, then accessing the console port places you in the system execution space. See [Chapter 3, “Enabling Multiple Context Mode,”](#) for more information about multiple context mode.



Note

If you want to use ASDM to configure the security appliance instead of the command-line interface, you can connect to the default management address of 192.168.1.1 (if your security appliance includes a factory default configuration. See the [“Factory Default Configurations”](#) section on page 2-1.). On the

ASA 5510 and higher adaptive security appliances, the interface to which you connect with ASDM is Management 0/0. For the ASA 5505 adaptive security appliance, the switch port to which you connect with ASDM is any port, except for Ethernet 0/0. For the PIX 515/515E security appliance, the interface to which you connect with ASDM is Ethernet 1. If you do not have a factory default configuration, follow the steps in this section to access the command-line interface. You can then configure the minimum parameters to access ASDM by entering the **setup** command.

To access the command-line interface, perform the following steps:

-
- Step 1** Connect a PC to the console port using the provided console cable, and connect to the console using a terminal emulator set for 9600 baud, 8 data bits, no parity, 1 stop bit, no flow control.
- See the hardware guide that came with your security appliance for more information about the console cable.
- Step 2** Press the **Enter** key to see the following prompt:
- ```
hostname>
```
- This prompt indicates that you are in user EXEC mode.
- Step 3** To access privileged EXEC mode, enter the following command:
- ```
hostname> enable
```
- The following prompt appears:
- ```
Password:
```
- Step 4** Enter the enable password at the prompt.
- By default, the password is blank, and you can press the **Enter** key to continue. See the [“Changing the Enable Password” section on page 8-1](#) to change the enable password.
- The prompt changes to:
- ```
hostname#
```
- To exit privileged mode, enter the **disable**, **exit**, or **quit** command.
- Step 5** To access global configuration mode, enter the following command:
- ```
hostname# configure terminal
```
- The prompt changes to the following:
- ```
hostname(config)#
```
- To exit global configuration mode, enter the **exit**, **quit**, or **end** command.
-

Setting Transparent or Routed Firewall Mode

You can set the security appliance to run in routed firewall mode (the default) or transparent firewall mode.

For multiple context mode, you can use only one firewall mode for all contexts. You must set the mode in the system execution space.

When you change modes, the security appliance clears the configuration because many commands are not supported for both modes. If you already have a populated configuration, be sure to back up your configuration before changing the mode; you can use this backup for reference when creating your new configuration. See the [“Backing Up Configuration Files” section on page 41-8](#). For multiple context mode, the system configuration is erased. This action removes any contexts from running. If you then re-add a context that has an existing configuration that was created for the wrong mode, the context configuration will not work correctly. Be sure to recreate your context configurations for the correct mode before you re-add them, or add new contexts with new paths for the new configurations.

If you download a text configuration to the security appliance that changes the mode with the **firewall transparent** command, be sure to put the command at the top of the configuration; the security appliance changes the mode as soon as it reads the command and then continues reading the configuration you downloaded. If the command is later in the configuration, the security appliance clears all the preceding lines in the configuration. See the [“Downloading Software or Configuration Files to Flash Memory” section on page 41-3](#) for information about downloading text files.

- To set the mode to transparent, enter the following command in the system execution space:

```
hostname(config)# firewall transparent
```

This command also appears in each context configuration for informational purposes only; you cannot enter this command in a context.

- To set the mode to routed, enter the following command in the system execution space:

```
hostname(config)# no firewall transparent
```

Working with the Configuration

This section describes how to work with the configuration. The security appliance loads the configuration from a text file, called the startup configuration. This file resides by default as a hidden file in internal Flash memory. You can, however, specify a different path for the startup configuration. (For more information, see [Chapter 41, “Managing Software, Licenses, and Configurations.”](#))

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

The information in this section applies to both single and multiple security contexts, except where noted. Additional information about contexts is in [Chapter 3, “Enabling Multiple Context Mode.”](#)

This section includes the following topics:

- [Saving Configuration Changes, page 2-6](#)
- [Copying the Startup Configuration to the Running Configuration, page 2-8](#)
- [Viewing the Configuration, page 2-8](#)
- [Clearing and Removing Configuration Settings, page 2-9](#)
- [Creating Text Configuration Files Offline, page 2-9](#)

Saving Configuration Changes

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

- [Saving Configuration Changes in Single Context Mode, page 2-7](#)

- [Saving Configuration Changes in Multiple Context Mode, page 2-7](#)

Saving Configuration Changes in Single Context Mode

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

```
hostname# write memory
```

**Note**

The **copy running-config startup-config** command is equivalent to the **write memory** command.

Saving Configuration Changes in Multiple Context Mode

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

- [Saving Each Context and System Separately, page 2-7](#)
- [Saving All Context Configurations at the Same Time, page 2-7](#)

Saving Each Context and System Separately

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

```
hostname# write memory
```

**Note**

The **copy running-config startup-config** command is equivalent to the **write memory** command.

For multiple context mode, context startup configurations can reside on external servers. In this case, the security appliance saves the configuration back to the server you identified in the context URL, except for an HTTP or HTTPS URL, which do not let you save the configuration to the server.

Saving All Context Configurations at the Same Time

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

```
hostname# write memory all [/noconfirm]
```

If you do not enter the **/noconfirm** keyword, you see the following prompt:

```
Are you sure [Y/N]:
```

After you enter **Y**, the security appliance saves the system configuration and each context. Context startup configurations can reside on external servers. In this case, the security appliance saves the configuration back to the server you identified in the context URL, except for an HTTP or HTTPS URL, which do not let you save the configuration to the server.

After the security appliance saves each context, the following message appears:

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

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

- For contexts that are not saved because of low memory, the following message appears:
The context 'context a' could not be saved due to Unavailability of resources

- For contexts that are not saved because the remote destination is unreachable, the following message appears:

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

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

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

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

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

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

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

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

Copying the Startup Configuration to the Running Configuration

Copy a new startup configuration to the running configuration using one of these options:

- To merge the startup configuration with the running configuration, enter the following command:

```
hostname(config)# copy startup-config running-config
```

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.

- To load the startup configuration and discard the running configuration, restart the security appliance by entering the following command:

```
hostname# reload
```

Alternatively, you can use the following commands to load the startup configuration and discard the running configuration without requiring a reboot:

```
hostname/contexta(config)# clear configure all  
hostname/contexta(config)# copy startup-config running-config
```

Viewing the Configuration

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

- To view the running configuration, enter the following command:

```
hostname# show running-config
```

- To view the running configuration of a specific command, enter the following command:

```
hostname# show running-config command
```

- To view the startup configuration, enter the following command:

```
hostname# show startup-config
```

Clearing and Removing Configuration Settings

To erase settings, enter one of the following commands.

- To clear all the configuration for a specified command, enter the following command:

```
hostname(config)# clear configure configurationcommand [level2configurationcommand]
```

This command clears all the current configuration for the specified configuration command. If you only want to clear the configuration for a specific version of the command, you can enter a value for *level2configurationcommand*.

For example, to clear the configuration for all **aaa** commands, enter the following command:

```
hostname(config)# clear configure aaa
```

To clear the configuration for only **aaa authentication** commands, enter the following command:

```
hostname(config)# clear configure aaa authentication
```

- To disable the specific parameters or options of a command, enter the following command:

```
hostname(config)# no configurationcommand [level2configurationcommand] qualifier
```

In this case, you use the **no** command to remove the specific configuration identified by *qualifier*.

For example, to remove a specific **nat** command, enter enough of the command to identify it uniquely as follows:

```
hostname(config)# no nat (inside) 1
```

- To erase the startup configuration, enter the following command:

```
hostname(config)# write erase
```

- To erase the running configuration, enter the following command:

```
hostname(config)# clear configure all
```

**Note**

In multiple context mode, if you enter **clear configure all** from the system configuration, you also remove all contexts and stop them from running.

Creating Text Configuration Files Offline

This guide describes how to use the CLI to configure the security appliance; when you save commands, the changes are written to a text file. Instead of using the CLI, however, you can edit a text file directly on your PC and paste a configuration at the configuration mode command-line prompt in its entirety, or line by line. Alternatively, you can download a text file to the security appliance internal Flash memory. See [Chapter 41, “Managing Software, Licenses, and Configurations,”](#) for information on downloading the configuration file to the security appliance.

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

```
hostname(config)# context a
```

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

```
context a
```

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



CHAPTER 3

Enabling Multiple Context Mode

This chapter describes how to use security contexts and enable multiple context mode. This chapter includes the following sections:

- [Security Context Overview, page 3-1](#)
- [Enabling or Disabling Multiple Context Mode, page 3-10](#)

Security Context Overview

You can partition a single security appliance into multiple virtual devices, known as security contexts. Each context is an independent device, with its own security policy, interfaces, and administrators. Multiple contexts are similar to having multiple standalone devices. Many features are supported in multiple context mode, including routing tables, firewall features, IPS, and management. Some features are not supported, including VPN and dynamic routing protocols.

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

- [Common Uses for Security Contexts, page 3-1](#)
- [Unsupported Features, page 3-2](#)
- [Context Configuration Files, page 3-2](#)
- [How the Security Appliance Classifies Packets, page 3-3](#)
- [Cascading Security Contexts, page 3-8](#)
- [Management Access to Security Contexts, page 3-9](#)

Common Uses for Security Contexts

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

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

Unsupported Features

Multiple context mode does not support the following features:

- Dynamic routing protocols

Security contexts support only static routes. You cannot enable OSPF or RIP in multiple context mode.

- VPN
- Multicast

Context Configuration Files

This section describes how the security appliance implements multiple context mode configurations and includes the following sections:

- [Context Configurations, page 3-2](#)
- [System Configuration, page 3-2](#)
- [Admin Context Configuration, page 3-2](#)

Context Configurations

The security appliance includes a configuration for each context that identifies the security policy, interfaces, and almost all the options you can configure on a standalone device. You can store context configurations on the internal Flash memory or the external Flash memory card, or you can download them from a TFTP, FTP, or HTTP(S) server.

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 security appliance. The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses one of the contexts that is designated as the *admin context*. The 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 Security Appliance Classifies Packets

Each packet that enters the security appliance must be classified, so that the security appliance can determine to which context to send a packet. This section includes the following topics:

- [Valid Classifier Criteria, page 3-3](#)
- [Invalid Classifier Criteria, page 3-4](#)
- [Classification Examples, page 3-5](#)

**Note**

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

Valid Classifier Criteria

This section describes the criteria used by the classifier, and includes the following topics:

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

Unique Interfaces

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

Unique MAC Addresses

If multiple contexts share an interface, then the classifier uses the interface MAC address. The security appliance lets you assign a different MAC address in each context to the same shared interface, whether it is a shared physical interface or a shared subinterface. By default, shared interfaces do not have unique MAC addresses; the interface uses the physical interface burned-in MAC address in every context. An upstream router cannot route directly to a context without unique MAC addresses. You can set the MAC addresses manually when you configure each interface (see the [“Configuring the Interface” section on page 7-2](#)), or you can automatically generate MAC addresses (see the [“Automatically Assigning MAC Addresses to Context Interfaces” section on page 6-11](#)).

NAT Configuration

If you do not have unique MAC addresses, then the classifier intercepts the packet and performs a destination IP address lookup. All other fields are ignored; only the destination IP address is used. To use the destination address for classification, the classifier must have knowledge about the subnets located behind each security context. The classifier relies on the NAT configuration to determine the subnets in each context. The classifier matches the destination IP address to either a **static** command or a **global** command. In the case of the **global** command, the classifier does not need a matching **nat** command or an active NAT session to classify the packet. Whether the packet can communicate with the destination IP address after classification depends on how you configure NAT and NAT control.

For example, the classifier gains knowledge about subnets 10.10.10.0, 10.20.10.0 and 10.30.10.0 when the context administrators configure **static** commands in each context:

- Context A:

```
static (inside,shared) 10.10.10.0 10.10.10.0 netmask 255.255.255.0
```

- Context B:

```
static (inside,shared) 10.20.10.0 10.20.10.0 netmask 255.255.255.0
```

- Context C:

```
static (inside,shared) 10.30.10.0 10.30.10.0 netmask 255.255.255.0
```

**Note**

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

Invalid Classifier Criteria


The following configurations are not used for packet classification:

- NAT exemption—The classifier does not use a NAT exemption configuration for classification purposes because NAT exemption does not identify a mapped interface.
- Routing table—If a context includes a static route that points to an external router as the next-hop to a subnet, and a different context includes a **static** command for the same subnet, then the classifier uses the **static** command to classify packets destined for that subnet and ignores the static route.

Classification Examples

Figure 3-2 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 3-1 Packet Classification with a Shared Interface using MAC Addresses



153367

Figure 3-2 shows multiple contexts sharing an outside interface without MAC addresses assigned. The classifier assigns the packet to Context B because Context B includes the address translation that matches the destination address.

Figure 3-2 Packet Classification with a Shared Interface using NAT




Note that all new incoming traffic must be classified, even from inside networks. Figure 3-3 shows a host on the Context B inside network accessing the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 0/1.3, which is assigned to Context B.




Note

If you share an *inside* interface and do not use unique MAC addresses, the classifier imposes some major restrictions. The classifier relies on the address translation configuration to classify the packet within a context, and you must translate the *destination* addresses of the traffic. Because you do not usually perform NAT on outside addresses, sending packets from inside to outside on a shared interface is not always possible; the outside network is large, (the Web, for example), and addresses are not predictable for an outside NAT configuration. If you share an inside interface, we suggest you use unique MAC addresses.

Figure 3-3 Incoming Traffic from Inside Networks

For transparent firewalls, you must use unique interfaces. Figure 3-4 shows a host on the Context B inside network accessing the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 1/0.3, which is assigned to Context B.

Figure 3-4 Transparent Firewall Contexts



Cascading Security Contexts

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




Note

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

Figure 3-5 shows a gateway context with two contexts behind the gateway.

Figure 3-5 Cascading Contexts



Management Access to Security Contexts

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

- [System Administrator Access, page 3-9](#)
- [Context Administrator Access, page 3-10](#)

System Administrator Access

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

- Access the security appliance console.
From the console, you access the system execution space.
- Access the admin context using Telnet, SSH, or ASDM.

See [Chapter 40, “Managing System Access,”](#) to enable Telnet, SSH, and SDM access.

As the system administrator, you can access all contexts.

When you change to a context from admin or the system, your username changes to the default “enable_15” username. If you configured command authorization in that context, you need to either configure authorization privileges for the “enable_15” user, or you can log in as a different name for which you provide sufficient privileges in the command authorization configuration for the context. To log in with a username, enter the **login** command. For example, you log in to the admin context with the

username “admin.” The admin context does not have any command authorization configuration, but all other contexts include command authorization. For convenience, each context configuration includes a user “admin” with maximum privileges. When you change from the admin context to context A, your username is altered, so you must log in again as “admin” by entering the **login** command. When you change to context B, you must again enter the **login** command to log in as “admin.”

The system execution space does not support any AAA commands, but you can configure its own enable password, as well as usernames in the local database to provide individual logins.

Context Administrator Access

You can access a context using Telnet, SSH, or ASDM. If you log in to a non-admin context, you can only access the configuration for that context. You can provide individual logins to the context. See [Chapter 40, “Managing System Access,”](#) to enable Telnet, SSH, and SDM access and to configure management authentication.

Enabling or Disabling Multiple Context Mode

Your security appliance might already be configured for multiple security contexts depending on how you ordered it from Cisco. If you are upgrading, however, you might need to convert from single mode to multiple mode by following the procedures in this section. ASDM does not support changing modes, so you need to change modes using the CLI.

This section includes the following topics:

- [Backing Up the Single Mode Configuration, page 3-10](#)
- [Enabling Multiple Context Mode, page 3-10](#)
- [Restoring Single Context Mode, page 3-11](#)

Backing Up the Single Mode Configuration

When you convert from single mode to multiple mode, the security appliance converts the running configuration into two files. The original startup configuration is not saved, so if it differs from the running configuration, you should back it up before proceeding.

Enabling Multiple Context Mode

The context mode (single or multiple) is not stored in the configuration file, even though it does endure reboots. If you need to copy your configuration to another device, set the mode on the new device to match using the **mode** command.

When you convert from single mode to multiple mode, the security appliance converts the running configuration into two files: a new startup configuration that comprises the system configuration, and admin.cfg that comprises the admin context (in the root directory of the internal Flash memory). The original running configuration is saved as old_running.cfg (in the root directory of the internal Flash memory). The original startup configuration is not saved. The security appliance automatically adds an entry for the admin context to the system configuration with the name “admin.”

To enable multiple mode, enter the following command:

```
hostname(config)# mode multiple
```


You are prompted to reboot the security appliance.

Restoring Single Context Mode

If you convert from multiple mode to single mode, you might want to first copy a full startup configuration (if available) to the security appliance; the system configuration inherited from multiple mode is not a complete functioning configuration for a single mode device. Because the system configuration does not have any network interfaces as part of its configuration, you must access the security appliance from the console to perform the copy.

To copy the old running configuration to the startup configuration and to change the mode to single mode, perform the following steps in the system execution space:

-
- Step 1** To copy the backup version of your original running configuration to the current startup configuration, enter the following command in the system execution space:

```
hostname(config)# copy flash:old_running.cfg startup-config
```

- Step 2** To set the mode to single mode, enter the following command in the system execution space:

```
hostname(config)# mode single
```

The security appliance reboots.



CHAPTER 4

Configuring Switch Ports and VLAN Interfaces for the Cisco ASA 5505 Adaptive Security Appliance

This chapter describes how to configure the switch ports and VLAN interfaces of the ASA 5505 adaptive security appliance.



Note

To configure interfaces of other models, see [Chapter 5, “Configuring Ethernet Settings and Subinterfaces,”](#) and [Chapter 7, “Configuring Interface Parameters.”](#)

This chapter includes the following sections:

- [Interface Overview, page 4-1](#)
- [Configuring VLAN Interfaces, page 4-5](#)
- [Configuring Switch Ports as Access Ports, page 4-9](#)
- [Configuring a Switch Port as a Trunk Port, page 4-11](#)
- [Allowing Communication Between VLAN Interfaces on the Same Security Level, page 4-13](#)

Interface Overview

This section describes the ports and interfaces of the ASA 5505 adaptive security appliance, and includes the following topics:

- [Understanding ASA 5505 Ports and Interfaces, page 4-2](#)
- [Maximum Active VLAN Interfaces for Your License, page 4-2](#)
- [Default Interface Configuration, page 4-4](#)
- [VLAN MAC Addresses, page 4-4](#)
- [Power Over Ethernet, page 4-4](#)
- [Security Level Overview, page 4-5](#)

Understanding ASA 5505 Ports and Interfaces

The ASA 5505 adaptive security appliance supports a built-in switch. There are two kinds of ports and interfaces that you need to configure:

- Physical switch ports—The adaptive security appliance has eight Fast Ethernet switch ports that forward traffic at Layer 2, using the switching function in hardware. Two of these ports are PoE ports. See the [“Power Over Ethernet” section on page 4-4](#) for more information. You can connect these interfaces directly to user equipment such as PCs, IP phones, or a DSL modem. Or you can connect to another switch.
- Logical VLAN interfaces—In routed mode, these interfaces forward traffic between VLAN networks at Layer 3, using the configured security policy to apply firewall and VPN services. In transparent mode, these interfaces forward traffic between the VLANs on the same network at Layer 2, using the configured security policy to apply firewall services. See the [“Maximum Active VLAN Interfaces for Your License”](#) section for more information about the maximum VLAN interfaces. VLAN interfaces let you divide your equipment into separate VLANs, for example, home, business, and Internet VLANs.

To segregate the switch ports into separate VLANs, you assign each switch port to a VLAN interface. Switch ports on the same VLAN can communicate with each other using hardware switching. But when a switch port on VLAN 1 wants to communicate with a switch port on VLAN 2, then the adaptive security appliance applies the security policy to the traffic and routes or bridges between the two VLANs.

**Note**

Subinterfaces are not available for the ASA 5505 adaptive security appliance.

Maximum Active VLAN Interfaces for Your License


In transparent firewall mode, you can configure two active VLANs in the Base license and three active VLANs in the Security Plus license, one of which must be for failover.

In routed mode, you can configure up to three active VLANs with the Base license, and up to 20 active VLANs with the Security Plus license.

An active VLAN is a VLAN with a **nameif** command configured.

With the Base license, the third VLAN can only be configured to initiate traffic to one other VLAN. See [Figure 4-1](#) for an example network where the Home VLAN can communicate with the Internet, but cannot initiate contact with Business.

Figure 4-1 *ASA 5505 Adaptive Security Appliance with Base License*



With the Security Plus license, you can configure 20 VLAN interfaces. You can configure trunk ports to accommodate multiple VLANs per port.




Note

The ASA 5505 adaptive security appliance supports Active/Standby failover, but not Stateful failover.

See [Figure 4-2](#) for an example network.

Figure 4-2 *ASA 5505 Adaptive Security Appliance with Security Plus License*



Default Interface Configuration

If your adaptive security appliance includes the default factory configuration, your interfaces are configured as follows:

- The outside interface (security level 0) is VLAN 2.
Ethernet0/0 is assigned to VLAN 2 and is enabled.
The VLAN 2 IP address is obtained from the DHCP server.
- The inside interface (security level 100) is VLAN 1.
Ethernet 0/1 through Ethernet 0/7 are assigned to VLAN 1 and is enabled.
VLAN 1 has IP address 192.168.1.1.

Restore the default factory configuration using the **configure factory-default** command.

Use the procedures in this chapter to modify the default configuration, for example, to add VLAN interfaces.

If you do not have a factory default configuration, all switch ports are in VLAN 1, but no other parameters are configured.

VLAN MAC Addresses

In routed firewall mode, all VLAN interfaces share a MAC address. Ensure that any connected switches can support this scenario. If the connected switches require unique MAC addresses, you can manually assign MAC addresses.

In transparent firewall mode, each VLAN has a unique MAC address. You can override the generated MAC addresses if desired by manually assigning MAC addresses.

Power Over Ethernet

Ethernet 0/6 and Ethernet 0/7 support PoE for devices such as IP phones or wireless access points. If you install a non-PoE device or do not connect to these switch ports, the adaptive security appliance does not supply power to the switch ports.

If you shut down the switch port using the **shutdown** command, you disable power to the device. Power is restored when you enter **no shutdown**. See the [“Configuring Switch Ports as Access Ports” section on page 4-9](#) for more information about shutting down a switch port.

To view the status of PoE switch ports, including the type of device connected (Cisco or IEEE 802.3af), use the **show power inline** command.

Monitoring Traffic Using SPAN

If you want to monitor traffic that enters or exits one or more switch ports, you can enable SPAN, also known as switch port monitoring. The port for which you enable SPAN (called the destination port) receives a copy of every packet transmitted or received on a specified source port. The SPAN feature lets you attach a sniffer to the destination port so you can monitor all traffic; without SPAN, you would have to attach a sniffer to every port you want to monitor. You can only enable SPAN for one destination port.

See the **switchport monitor** command in the *Cisco Security Appliance Command Reference* for more information.

Security Level Overview

Each VLAN interface must have a security level in the range 0 to 100 (from lowest to highest). For example, you should assign your most secure network, such as the inside business network, to level 100. The outside network connected to the Internet can be level 0. Other networks, such as a home network 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 access list 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. See the [“Allowing Communication Between VLAN Interfaces on the Same Security Level”](#) section on page 4-13 for more information.
- 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 adaptive security appliance.
- Filtering—HTTP(S) and FTP filtering applies only for outbound connections (from a higher level to a lower level).

For same security interfaces, you can filter traffic in either direction.

- NAT control—When you enable NAT control, you must configure NAT for hosts on a higher security interface (inside) when they access hosts on a lower security interface (outside).

Without NAT control, or for same security interfaces, you can choose to use NAT between any interface, or you can choose not to use NAT. Keep in mind that configuring NAT for an outside interface might require a special keyword.

- **established** command—This command allows return connections from a lower security host to a higher security host if there is already an established connection from the higher level host to the lower level host.

For same security interfaces, you can configure **established** commands for both directions.

Configuring VLAN Interfaces

For each VLAN to pass traffic, you need to configure an interface name (the **nameif** command), and for routed mode, an IP address. You should also change the security level from the default, which is 0. If you name an interface “inside” and you do not set the security level explicitly, then the adaptive security appliance sets the security level to 100.

For information about how many VLANs you can configure, see the [“Maximum Active VLAN Interfaces for Your License”](#) section on page 4-2.

**Note**

If you are using failover, do not use this procedure to name interfaces that you are reserving for failover communications. See [Chapter 14, “Configuring Failover,”](#) to configure the failover link.

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.

To configure a VLAN interface, perform the following steps:

- Step 1** To specify the VLAN ID, enter the following command:

```
hostname(config)# interface vlan number
```

Where the *number* is between 1 and 4090.

For example, enter the following command:

```
hostname(config)# interface vlan 100
```

To remove this VLAN interface and all associated configuration, enter the **no interface vlan** command. Because this interface also includes the interface name configuration, and the name is used in other commands, those commands are also removed.

- Step 2** (Optional) For the Base license, allow this interface to be the third VLAN by limiting it from initiating contact to one other VLAN using the following command:

```
hostname(config-if)# no forward interface vlan number
```

Where *number* specifies the VLAN ID to which this VLAN interface cannot initiate traffic.

With the Base license, you can only configure a third VLAN if you use this command to limit it.

For example, you have one VLAN assigned to the outside for Internet access, one VLAN assigned to an inside business network, and a third VLAN assigned to your home network. The home network does not need to access the business network, so you can use the **no forward interface** command on the home VLAN; the business network can access the home network, but the home network cannot access the business network.

If you already have two VLAN interfaces configured with a **nameif** command, be sure to enter the **no forward interface** command before the **nameif** command on the third interface; the adaptive security appliance does not allow three fully functioning VLAN interfaces with the Base license on the ASA 5505 adaptive security appliance.

**Note**

If you upgrade to the Security Plus license, you can remove this command and achieve full functionality for this interface. If you leave this command in place, this interface continues to be limited even after upgrading.

- Step 3** To name the interface, enter the following command:

```
hostname(config-if)# nameif 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. Do not enter the **no** form, because that command causes all commands that refer to that name to be deleted.

- Step 4** To set the security level, enter the following command:

```
hostname(config-if)# security-level number
```


Where *number* is an integer between 0 (lowest) and 100 (highest).

Step 5 (Routed mode only) To set the IP address, enter one of the following commands.



Note To set an IPv6 address, see the [“Configuring IPv6 on an Interface”](#) section on page 12-3.

To set the management IP address for transparent firewall mode, see the [“Setting the Management IP Address for a Transparent Firewall”](#) section on page 8-5. In transparent mode, you do not set the IP address for each interface, but rather for the whole adaptive security appliance or context.

For failover, you must set the IP address an standby address manually; DHCP and PPPoE are not supported.

- To set the IP address manually, enter the following command:

```
hostname(config-if)# ip address ip_address [mask] [standby ip_address]
```

The **standby** keyword and address is used for failover. See [Chapter 14, “Configuring Failover,”](#) for more information.

- To obtain an IP address from a DHCP server, enter the following command:

```
hostname(config-if)# ip address dhcp [setroute]
```

Reenter this command to reset the DHCP lease and request a new lease.

If you do not enable the interface using the **no shutdown** command before you enter the **ip address dhcp** command, some DHCP requests might not be sent.

- To obtain an IP address from a PPPoE server, see [Chapter 35, “Configuring the PPPoE Client.”](#)

Step 6 (Optional) To assign a private MAC address to this interface, enter the following command:

```
hostname(config-if)# mac-address mac_address [standby mac_address]
```

By default in routed mode, all VLANs use the same MAC address. In transparent mode, the VLANs use unique MAC addresses. You might want to set unique VLANs or change the generated VLANs if your switch requires it, or for access control purposes.

Step 7 (Optional) To set an interface to management-only mode, so that it does not allow through traffic, enter the following command:

```
hostname(config-if)# management-only
```

Step 8 By default, VLAN interfaces are enabled. To enable the interface, if it is not already enabled, enter the following command:

```
hostname(config-if)# no shutdown
```

To disable the interface, enter the **shutdown** command.

The following example configures seven VLAN interfaces, including the failover interface which is configured separately using the **failover lan** command:

```
hostname(config)# interface vlan 100
hostname(config-if)# nameif outside
hostname(config-if)# security-level 0
hostname(config-if)# ip address 10.1.1.1 255.255.255.0
```

```

hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 200
hostname(config-if)# nameif inside
hostname(config-if)# security-level 100
hostname(config-if)# ip address 10.2.1.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 201
hostname(config-if)# nameif dept1
hostname(config-if)# security-level 90
hostname(config-if)# ip address 10.2.2.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 202
hostname(config-if)# nameif dept2
hostname(config-if)# security-level 90
hostname(config-if)# ip address 10.2.3.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 300
hostname(config-if)# nameif dmz
hostname(config-if)# security-level 50
hostname(config-if)# ip address 10.3.1.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 400
hostname(config-if)# nameif backup-isp
hostname(config-if)# security-level 50
hostname(config-if)# ip address 10.1.2.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# failover lan faillink vlan500
hostname(config)# failover interface ip faillink 10.4.1.1 255.255.255.0 standby 10.4.1.2
255.255.255.0

```

The following example configures three VLAN interfaces for the Base license. The third home interface cannot forward traffic to the business interface.

```

hostname(config)# interface vlan 100
hostname(config-if)# nameif outside
hostname(config-if)# security-level 0
hostname(config-if)# ip address dhcp
hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 200
hostname(config-if)# nameif business
hostname(config-if)# security-level 100
hostname(config-if)# ip address 10.1.1.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 300
hostname(config-if)# no forward interface vlan 200
hostname(config-if)# nameif home
hostname(config-if)# security-level 50
hostname(config-if)# ip address 10.2.1.1 255.255.255.0
hostname(config-if)# no shutdown

```

Configuring Switch Ports as Access Ports

By default, all switch ports are shut down. To assign a switch port to one VLAN, configure it as an access port. To create a trunk port to carry multiple VLANs, see the [“Configuring a Switch Port as a Trunk Port” section on page 4-11](#).

By default, the speed and duplex for switch ports are set to auto-negotiate. 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.

**Caution**

The ASA 5505 adaptive security appliance does not support Spanning Tree Protocol for loop detection in the network. Therefore you must ensure that any connection with the adaptive security appliance does not end up in a network loop.

To configure a switch port, perform the following steps:

- Step 1** To specify the switch port you want to configure, enter the following command:

```
hostname(config)# interface ethernet0/port
```

Where *port* is 0 through 7. For example, enter the following command:

```
hostname(config)# interface ethernet0/1
```

- Step 2** To assign this switch port to a VLAN, enter the following command:

```
hostname(config-if)# switchport access vlan number
```

Where *number* is the VLAN ID, between 1 and 4090.

**Note**

You might assign multiple switch ports to the primary or backup VLANs if the Internet access device includes Layer 2 redundancy.

- Step 3** (Optional) To prevent the switch port from communicating with other protected switch ports on the same VLAN, enter the following command:

```
hostname(config-if)# switchport protected
```

You might want to prevent switch ports from communicating with each other if the devices on those switch ports are primarily accessed from other VLANs, you do not need to allow intra-VLAN access, and you want to isolate the devices from each other in case of infection or other security breach. For example, if you have a DMZ that hosts three web servers, you can isolate the web servers from each other if you apply the **switchport protected** command to each switch port. The inside and outside networks can both communicate with all three web servers, and vice versa, but the web servers cannot communicate with each other.

- Step 4** (Optional) To set the speed, enter the following command:

```
hostname(config-if)# speed {auto | 10 | 100}
```

The **auto** setting is the default. If you set the speed to anything other than **auto** on PoE ports Ethernet 0/6 or 0/7, then Cisco IP phones and Cisco wireless access points that do not support IEEE 802.3af will not be detected and supplied with power.

Step 5 (Optional) To set the duplex, enter the following command:

```
hostname(config-if)# duplex {auto | full | half}
```

The **auto** setting is the default. If you set the duplex to anything other than **auto** on PoE ports Ethernet 0/6 or 0/7, then Cisco IP phones and Cisco wireless access points that do not support IEEE 802.3af will not be detected and supplied with power.

Step 6 To enable the switch port, if it is not already enabled, enter the following command:

```
hostname(config-if)# no shutdown
```

To disable the switch port, enter the **shutdown** command.

The following example configures five VLAN interfaces, including the failover interface which is configured using the **failover lan** command:

```
hostname(config)# interface vlan 100
hostname(config-if)# nameif outside
hostname(config-if)# security-level 0
hostname(config-if)# ip address 10.1.1.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 200
hostname(config-if)# nameif inside
hostname(config-if)# security-level 100
hostname(config-if)# ip address 10.2.1.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 300
hostname(config-if)# nameif dmz
hostname(config-if)# security-level 50
hostname(config-if)# ip address 10.3.1.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 400
hostname(config-if)# nameif backup-isp
hostname(config-if)# security-level 50
hostname(config-if)# ip address 10.1.2.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# failover lan faillink vlan500
hostname(config)# failover interface ip faillink 10.4.1.1 255.255.255.0 standby 10.4.1.2
255.255.255.0

hostname(config)# interface ethernet 0/0
hostname(config-if)# switchport access vlan 100
hostname(config-if)# no shutdown

hostname(config-if)# interface ethernet 0/1
hostname(config-if)# switchport access vlan 200
hostname(config-if)# no shutdown

hostname(config-if)# interface ethernet 0/2
hostname(config-if)# switchport access vlan 300
hostname(config-if)# no shutdown

hostname(config-if)# interface ethernet 0/3
```

```
hostname(config-if)# switchport access vlan 400
hostname(config-if)# no shutdown

hostname(config-if)# interface ethernet 0/4
hostname(config-if)# switchport access vlan 500
hostname(config-if)# no shutdown
```

Configuring a Switch Port as a Trunk Port

By default, all switch ports are shut down. This procedure tells how to create a trunk port that can carry multiple VLANs using 802.1Q tagging. Trunk mode is available only with the Security Plus license.

To create an access port, where an interface is assigned to only one VLAN, see the [“Configuring Switch Ports as Access Ports” section on page 4-9](#).

By default, the speed and duplex for switch ports are set to auto-negotiate. 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.

To configure a trunk port, perform the following steps:

-
- Step 1** To specify the switch port you want to configure, enter the following command:

```
hostname(config)# interface ethernet0/port
```

Where *port* is 0 through 7. For example, enter the following command:

```
hostname(config)# interface ethernet0/1
```

- Step 2** To assign VLANs to this trunk, enter one or more of the following commands.

- To assign native VLANs, enter the following command:

```
hostname(config-if)# switchport trunk native vlan vlan_id
```

where the *vlan_id* is a single VLAN ID between 1 and 4090.

Packets on the native VLAN are not modified when sent over the trunk. For example, if a port has VLANs 2, 3 and 4 assigned to it, and VLAN 2 is the native VLAN, then packets on VLAN 2 that egress the port are not modified with an 802.1Q header. Frames which ingress (enter) this port and have no 802.1Q header are put into VLAN 2.

Each port can only have one native VLAN, but every port can have either the same or a different native VLAN.

- To assign VLANs, enter the following command:

```
hostname(config-if)# switchport trunk allowed vlan vlan_range
```

where the *vlan_range* (with VLANs between 1 and 4090) can be identified in one of the following ways:

A single number (n)

A range (n-x)

Separate numbers and ranges by commas, for example:

5,7-10,13,45-100

You can enter spaces instead of commas, but the command is saved to the configuration with commas.

You can include the native VLAN in this command, but it is not required; the native VLAN is passed whether it is included in this command or not.

This switch port cannot pass traffic until you assign at least one VLAN to it, native or non-native.

- Step 3** To make this switch port a trunk port, enter the following command:

```
hostname(config-if)# switchport mode trunk
```

To restore this port to access mode, enter the **switchport mode access** command.

- Step 4** (Optional) To prevent the switch port from communicating with other protected switch ports on the same VLAN, enter the following command:

```
hostname(config-if)# switchport protected
```

You might want to prevent switch ports from communicating with each other if the devices on those switch ports are primarily accessed from other VLANs, you do not need to allow intra-VLAN access, and you want to isolate the devices from each other in case of infection or other security breach. For example, if you have a DMZ that hosts three web servers, you can isolate the web servers from each other if you apply the **switchport protected** command to each switch port. The inside and outside networks can both communicate with all three web servers, and vice versa, but the web servers cannot communicate with each other.

- Step 5** (Optional) To set the speed, enter the following command:

```
hostname(config-if)# speed {auto | 10 | 100}
```

The **auto** setting is the default.

- Step 6** (Optional) To set the duplex, enter the following command:

```
hostname(config-if)# duplex {auto | full | half}
```

The **auto** setting is the default.

- Step 7** To enable the switch port, if it is not already enabled, enter the following command:

```
hostname(config-if)# no shutdown
```

To disable the switch port, enter the **shutdown** command.

The following example configures seven VLAN interfaces, including the failover interface which is configured using the **failover lan** command. VLANs 200, 201, and 202 are trunked on Ethernet 0/1.

```
hostname(config)# interface vlan 100
hostname(config-if)# nameif outside
hostname(config-if)# security-level 0
hostname(config-if)# ip address 10.1.1.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 200
hostname(config-if)# nameif inside
hostname(config-if)# security-level 100
hostname(config-if)# ip address 10.2.1.1 255.255.255.0
hostname(config-if)# no shutdown

hostname(config-if)# interface vlan 201
hostname(config-if)# nameif dept1
```

```
hostname(config-if) # security-level 90
hostname(config-if) # ip address 10.2.2.1 255.255.255.0
hostname(config-if) # no shutdown

hostname(config-if) # interface vlan 202
hostname(config-if) # nameif dept2
hostname(config-if) # security-level 90
hostname(config-if) # ip address 10.2.3.1 255.255.255.0
hostname(config-if) # no shutdown

hostname(config-if) # interface vlan 300
hostname(config-if) # nameif dmz
hostname(config-if) # security-level 50
hostname(config-if) # ip address 10.3.1.1 255.255.255.0
hostname(config-if) # no shutdown

hostname(config-if) # interface vlan 400
hostname(config-if) # nameif backup-isp
hostname(config-if) # security-level 50
hostname(config-if) # ip address 10.1.2.1 255.255.255.0
hostname(config-if) # no shutdown

hostname(config-if) # failover lan faillink vlan500
hostname(config) # failover interface ip faillink 10.4.1.1 255.255.255.0 standby 10.4.1.2
255.255.255.0

hostname(config) # interface ethernet 0/0
hostname(config-if) # switchport access vlan 100
hostname(config-if) # no shutdown

hostname(config-if) # interface ethernet 0/1
hostname(config-if) # switchport mode trunk
hostname(config-if) # switchport trunk allowed vlan 200-202
hostname(config-if) # switchport trunk native vlan 5
hostname(config-if) # no shutdown

hostname(config-if) # interface ethernet 0/2
hostname(config-if) # switchport access vlan 300
hostname(config-if) # no shutdown

hostname(config-if) # interface ethernet 0/3
hostname(config-if) # switchport access vlan 400
hostname(config-if) # no shutdown

hostname(config-if) # interface ethernet 0/4
hostname(config-if) # switchport access vlan 500
hostname(config-if) # no shutdown
```

Allowing Communication Between VLAN Interfaces on the Same Security Level

By default, interfaces on the same security level cannot communicate with each other. Allowing communication between same security interfaces lets traffic flow freely between all same security interfaces without access lists.

**Note**

If you enable NAT control, you do not need to configure NAT between same security level interfaces. See the [“NAT and Same Security Level Interfaces” section on page 17-13](#) for more information on NAT and same security level interfaces.

If you enable same security interface communication, you can still configure interfaces at different security levels as usual.

To enable interfaces on the same security level so that they can communicate with each other, enter the following command:

```
hostname(config)# same-security-traffic permit inter-interface
```

To disable this setting, use the **no** form of this command.



CHAPTER 5

Configuring Ethernet Settings and Subinterfaces

This chapter describes how to configure and enable physical Ethernet interfaces and how to add subinterfaces. If you have both fiber and copper Ethernet ports (for example, on the 4GE SSM for the ASA 5510 and higher series adaptive security appliance), this chapter describes how to configure the interface media type.

In single context mode, complete the procedures in this chapter and then continue your interface configuration in [Chapter 7, “Configuring Interface Parameters.”](#) In multiple context mode, complete the procedures in this chapter in the system execution space, then assign interfaces and subinterfaces to contexts according to [Chapter 6, “Adding and Managing Security Contexts,”](#) and finally configure the interface parameters within each context according to [Chapter 7, “Configuring Interface Parameters.”](#)



Note

To configure interfaces for the ASA 5505 adaptive security appliance, see [Chapter 4, “Configuring Switch Ports and VLAN Interfaces for the Cisco ASA 5505 Adaptive Security Appliance.”](#)

This chapter includes the following sections:

- [Configuring and Enabling RJ-45 Interfaces, page 5-1](#)
- [Configuring and Enabling Fiber Interfaces, page 5-3](#)
- [Configuring and Enabling VLAN Subinterfaces and 802.1Q Trunking, page 5-3](#)

Configuring and Enabling RJ-45 Interfaces

This section describes how to configure Ethernet settings for physical interfaces, and how to enable the interface. By default, all physical interfaces are shut down. You must enable the physical interface before any traffic can pass through it or through a subinterface. For multiple context mode, if you allocate a physical interface or subinterface to a context, the interfaces are enabled by default in the context. However, before traffic can pass through the context interface, you must also enable the interface in the system configuration according to this procedure.

By default, the speed and duplex for copper (RJ-45) interfaces are set to auto-negotiate.

The ASA 5550 adaptive security appliance and the 4GE SSM for the ASA 5510 and higher adaptive security appliance includes two connector types: copper RJ-45 and fiber SFP. RJ-45 is the default. If you want to configure the security appliance to use the fiber SFP connectors, see the [“Configuring and Enabling Fiber Interfaces”](#) section on page 5-3.

For RJ-45 interfaces on the ASA 5500 series adaptive security appliance, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation

phase. Either the speed or duplex must be set to auto-negotiate to enable Auto-MDI/MDIX for the interface. If you explicitly set both the speed and duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled. For Gigabit Ethernet, when the speed and duplex are set to 1000 and full, then the interface always auto-negotiates; therefore Auto-MDI/MDIX is always enabled and you cannot disable it.

To enable the interface, or to set a specific speed and duplex, perform the following steps:

Step 1 To specify the interface you want to configure, enter the following command:

```
hostname(config)# interface physical_interface
```

The *physical_interface* ID includes the type, slot, and port number as *type[slot/port]*.

The physical interface types include the following:

- **ethernet**
- **gigabitethernet**

For the PIX 500 series security appliance, enter the type followed by the port number, for example, **ethernet0**.

For the ASA 5500 series adaptive security appliance, enter the type followed by slot/port, for example, **gigabitethernet0/1**. Interfaces that are built into the chassis are assigned to slot 0, while interfaces on the 4GE SSM are assigned to slot 1.

The ASA 5500 series adaptive security appliance also includes the following type:

- **management**

The management interface is a Fast Ethernet interface designed for management traffic only, and is specified as **management0/0**. You can, however, use it for through traffic if desired (see the **management-only** command). In transparent firewall mode, you can use the management interface in addition to the two interfaces allowed for through traffic. You can also add subinterfaces to the management interface to provide management in each security context for multiple context mode.

Step 2 (Optional) To set the speed, enter the following command:

```
hostname(config-if)# speed {auto | 10 | 100 | 1000 | nonegotiate}
```

The **auto** setting is the default. The **speed nonegotiate** command disables link negotiation.

Step 3 (Optional) To set the duplex, enter the following command:

```
hostname(config-if)# duplex {auto | full | half}
```

The **auto** setting is the default.

Step 4 To enable the interface, enter the following command:

```
hostname(config-if)# no shutdown
```

To disable the interface, enter the **shutdown** command. If you enter the **shutdown** command for a physical interface, you also shut down all subinterfaces. If you shut down an interface in the system execution space, then that interface is shut down in all contexts that share it.

Configuring and Enabling Fiber Interfaces

This section describes how to configure Ethernet settings for physical interfaces, and how to enable the interface. By default, all physical interfaces are shut down. You must enable the physical interface before any traffic can pass through it or through a subinterface. For multiple context mode, if you allocate a physical interface or subinterface to a context, the interfaces are enabled by default in the context. However, before traffic can pass through the context interface, you must also enable the interface in the system configuration according to this procedure.

By default, the connectors used on the 4GE SSM or for built-in interfaces in slot 1 on the ASA 5550 adaptive security appliance are the RJ-45 connectors. To use the fiber SFP connectors, you must set the media type to SFP. The fiber interface has a fixed speed and does not support duplex, but you can set the interface to negotiate link parameters (the default) or not to negotiate.

To enable the interface, set the media type, or to set negotiation settings, perform the following steps:

-
- Step 1** To specify the interface you want to configure, enter the following command:

```
hostname(config)# interface gigabitethernet 1/port
```

The 4GE SSM interfaces are assigned to slot 1, as shown in the interface ID in the syntax (the interfaces built into the chassis are assigned to slot 0).

- Step 2** To set the media type to SFP, enter the following command:

```
hostname(config-if)# media-type sfp
```

To restore the default RJ-45, enter the **media-type rj45** command.

- Step 3** (Optional) To disable link negotiation, enter the following command:

```
hostname(config-if)# speed nonegotiate
```

For fiber Gigabit Ethernet interfaces, the default is **no speed nonegotiate**, which sets the speed to 1000 Mbps and enables link negotiation for flow-control parameters and remote fault information. The **speed nonegotiate** command disables link negotiation.

- Step 4** To enable the interface, enter the following command:

```
hostname(config-if)# no shutdown
```

To disable the interface, enter the **shutdown** command. If you enter the **shutdown** command for a physical interface, you also shut down all subinterfaces. If you shut down an interface in the system execution space, then that interface is shut down in all contexts that share it.

Configuring and Enabling VLAN Subinterfaces and 802.1Q Trunking

This section describes how to configure and enable a VLAN subinterface. An interface with one or more VLAN subinterfaces is automatically configured as an 802.1Q trunk.

You must enable the physical interface before any traffic can pass through an enabled subinterface (see the “[Configuring and Enabling RJ-45 Interfaces](#)” section on page 5-1 or the “[Configuring and Enabling Fiber Interfaces](#)” section on page 5-3). For multiple context mode, if you allocate a subinterface to a context, the interfaces are enabled by default in the context. However, before traffic can pass through the context interface, you must also enable the interface in the system configuration with this procedure.

Subinterfaces let you divide a physical interface into multiple logical interfaces that are tagged with different VLAN IDs. Because VLANs allow you to keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces or security appliances. This feature is particularly useful in multiple context mode so you can assign unique interfaces to each context.

To determine how many subinterfaces are allowed for your platform, see [Appendix A, “Feature Licenses and Specifications.”](#)



Note

If you use subinterfaces, you typically do not also want the physical interface to pass traffic, because the physical interface passes untagged packets. Because the physical interface must be enabled for the subinterface to pass traffic, ensure that the physical interface does not pass traffic by leaving out the **nameif** command. If you want to let the physical interface pass untagged packets, you can configure the **nameif** command as usual. See the “[Configuring Interface Parameters](#)” section on page 7-1 for more information about completing the interface configuration.

To add a subinterface and assign a VLAN to it, perform the following steps:

Step 1 To specify the new subinterface, enter the following command:

```
hostname(config)# interface physical_interface.subinterface
```

See the “[Configuring and Enabling RJ-45 Interfaces](#)” section for a description of the physical interface ID.

The *subinterface* ID is an integer between 1 and 4294967293.

For example, enter the following command:

```
hostname(config)# interface gigabitethernet0/1.100
```

Step 2 To specify the VLAN for the subinterface, enter the following command:

```
hostname(config-subif)# vlan vlan_id
```

The *vlan_id* is an integer between 1 and 4094. Some VLAN IDs might be reserved on connected switches, so check the switch documentation for more information.

You can only assign a single VLAN to a subinterface, and not to the physical interface. Each subinterface must have a VLAN ID before it can pass traffic. To change a VLAN ID, you do not need to remove the old VLAN ID with the **no** option; you can enter the **vlan** command with a different VLAN ID, and the security appliance changes the old ID.

Step 3 To enable the subinterface, enter the following command:

```
hostname(config-subif)# no shutdown
```

To disable the interface, enter the **shutdown** command. If you shut down an interface in the system execution space, then that interface is shut down in all contexts that share it.



CHAPTER 6

Adding and Managing Security Contexts

This chapter describes how to configure multiple security contexts on the security appliance, and includes the following sections:

- [Configuring Resource Management, page 6-1](#)
- [Configuring a Security Context, page 6-7](#)
- [Automatically Assigning MAC Addresses to Context Interfaces, page 6-11](#)
- [Changing Between Contexts and the System Execution Space, page 6-11](#)
- [Managing Security Contexts, page 6-12](#)

For information about how contexts work and how to enable multiple context mode, see [Chapter 3, “Enabling Multiple Context Mode.”](#)

Configuring Resource Management

By default, all security contexts have unlimited access to the resources of the security appliance, except where maximum limits per context are enforced. However, if you find that one or more contexts use too many resources, and they cause other contexts to be denied connections, for example, then you can configure resource management to limit the use of resources per context.

This section includes the following topics:

- [Classes and Class Members Overview, page 6-1](#)
- [Configuring a Class, page 6-4](#)

Classes and Class Members Overview

The security appliance manages resources by assigning contexts to resource classes. Each context uses the resource limits set by the class. This section includes the following topics:

- [Resource Limits, page 6-2](#)
- [Default Class, page 6-3](#)
- [Class Members, page 6-4](#)


Resource Limits

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

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


You can oversubscribe the security appliance by assigning more than 100 percent of a resource across all contexts. For example, you can set the Bronze class to limit connections to 20 percent per context, and then assign 10 contexts to the class for a total of 200 percent. If contexts concurrently use more than the system limit, then each context gets less than the 20 percent you intended. (See [Figure 6-1](#).)

Figure 6-1 Resource Oversubscription



If you assign an absolute value to a resource across all contexts that exceeds the practical limit of the security appliance, then the performance of the security appliance might be impaired.

The security appliance lets you assign unlimited access to one or more resources in a class, instead of a percentage or absolute number. When a resource is unlimited, contexts can use as much of the resource as the system has available or that is practically available. For example, Context A, B, and C are in the Silver Class, which limits each class member to 1 percent of the connections, for a total of 3 percent; but the three contexts are currently only using 2 percent combined. Gold Class has unlimited access to connections. The contexts in the Gold Class can use more than the 97 percent of “unassigned” connections; they can also use the 1 percent of connections not currently in use by Context A, B, and C, even if that means that Context A, B, and C are unable to reach their 3 percent combined limit. (See [Figure 6-2](#).) Setting unlimited access is similar to oversubscribing the security appliance, except that you have less control over how much you oversubscribe the system.

Figure 6-2 *Unlimited Resources*

153211

Default Class

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


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

By default, the default class provides unlimited access to resources for all contexts, except for the following limits, which are by default set to the maximum allowed per context:

- Telnet sessions—5 sessions.
- SSH sessions—5 sessions.
- IPSec sessions—5 sessions.
- MAC addresses—65,535 entries.

Figure 6-3 shows the relationship between the default class and other classes. Contexts A and C belong to classes with some limits set; other limits are inherited from the default class. Context B inherits no limits from default because all limits are set in its class, the Gold class. Context D was not assigned to a class, and is by default a member of the default class.

Figure 6-3 Resource Classes



104689

Class Members

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

Configuring a Class

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.

- Step 1** To specify the class name and enter the class configuration mode, enter the following command in the system execution space:

```
hostname(config)# class name
```

The *name* is a string up to 20 characters long. To set the limits for the default class, enter **default** for the name.

- Step 2** To set the resource limits, see the following options:

- To set all resource limits (shown in Table 6-1) to be unlimited, enter the following command:

```
hostname(config-resgmt)# limit-resource all 0
```


For example, you might want to create a class that includes the admin context that has no limitations. The default class has all resources set to unlimited by default.

- To set a particular resource limit, enter the following command:

```
hostname(config-resgmt)# limit-resource [rate] resource_name number[%]
```

For this particular resource, the limit overrides the limit set for **all**. Enter the **rate** argument to set the rate per second for certain resources. For resources that do not have a system limit, you cannot set the percentage (%) between 1 and 100; you can only set an absolute value. See [Table 6-1](#) for resources for which you can set the rate per second and which to not have a system limit.

[Table 6-1](#) lists the resource types and the limits. See also the **show resource types** command.

Table 6-1 **Resource Names and Limits**

Resource Name	Rate or Concurrent	Minimum and Maximum Number per Context	System Limit ¹	Description
mac-addresses	Concurrent	N/A	65,535	For transparent firewall mode, the number of MAC addresses allowed in the MAC address table.
conns	Concurrent or Rate	N/A	Concurrent connections: See the “Supported Platforms and Feature Licenses” section on page A-1 for the connection limit for your platform. Rate: N/A	TCP or UDP connections between any two hosts, including connections between one host and multiple other hosts.
inspects	Rate	N/A	N/A	Application inspections.
hosts	Concurrent	N/A	N/A	Hosts that can connect through the security appliance.
asdm	Concurrent	1 minimum 5 maximum	32	ASDM management sessions. Note ASDM sessions use two HTTPS connections: one for monitoring that is always present, and one for making configuration changes that is present only when you make changes. For example, the system limit of 32 ASDM sessions represents a limit of 64 HTTPS sessions.
ssh	Concurrent	1 minimum 5 maximum	100	SSH sessions.
syslogs	Rate	N/A	N/A	System log messages.
telnet	Concurrent	1 minimum 5 maximum	100	Telnet sessions.
xlates	Concurrent	N/A	N/A	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.

For example, to set the default class limit for conns to 10 percent instead of unlimited, enter the following commands:

```
hostname(config)# class default
hostname(config-class)# limit-resource conns 10%
```

All other resources remain at unlimited.

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

```
hostname(config)# class gold
```

```

hostname(config-class)# limit-resource mac-addresses 10000
hostname(config-class)# limit-resource conns 15%
hostname(config-class)# limit-resource rate conns 1000
hostname(config-class)# limit-resource rate inspects 500
hostname(config-class)# limit-resource hosts 9000
hostname(config-class)# limit-resource asdm 5
hostname(config-class)# limit-resource ssh 5
hostname(config-class)# limit-resource rate syslogs 5000
hostname(config-class)# limit-resource telnet 5
hostname(config-class)# limit-resource xlates 36000

```

Configuring a Security Context

The security context definition in the system configuration identifies the context name, configuration file URL, and interfaces that a context can use.



Note

If you do not have an admin context (for example, if you clear the configuration) then you must first specify the admin context name by entering the following command:

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

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

To add or change a context in the system configuration, perform the following steps:

- Step 1** To add or modify a context, enter the following command in the system execution space:

```
hostname(config)# context name
```

The *name* is a string up to 32 characters long. This name is case sensitive, so you can have two contexts named “customerA” and “CustomerA,” for example. You can use letters, digits, or hyphens, but you cannot start or end the name with a hyphen.

“System” or “Null” (in upper or lower case letters) are reserved names, and cannot be used.

- Step 2** (Optional) To add a description for this context, enter the following command:

```
hostname(config-ctx)# description text
```

- Step 3** To specify the interfaces you can use in the context, enter the command appropriate for a physical interface or for one or more subinterfaces.

- To allocate a physical interface, enter the following command:

```
hostname(config-ctx)# allocate-interface physical_interface [map_name]
[visible | invisible]
```

- To allocate one or more subinterfaces, enter the following command:

```
hostname(config-ctx)# allocate-interface
physical_interface.subinterface[-physical_interface.subinterface]
[map_name[-map_name]] [visible | invisible]
```

You can enter these commands multiple times to specify different ranges. If you remove an allocation with the **no** form of this command, then any context commands that include this interface are removed from the running configuration.

Transparent firewall mode allows only two interfaces to pass through traffic; however, on the ASA adaptive security appliance, you can use the dedicated management interface, Management 0/0, (either the physical interface or a subinterface) as a third interface for management traffic.

**Note**

The management interface for transparent mode does not flood a packet out the interface when that packet is not in the MAC address table.

You can assign the same interfaces to multiple contexts in routed mode, if desired. Transparent mode does not allow shared interfaces.

The *map_name* is an alphanumeric alias for the interface that can be used within the context instead of the interface ID. If you do not specify a mapped name, the interface ID is used within the context. For security purposes, you might not want the context administrator to know which interfaces are being used by the context.

A mapped name must start with a letter, end with a letter or digit, and have as interior characters only letters, digits, or an underscore. For example, you can use the following names:

int0

inta

int_0

For subinterfaces, you can specify a range of mapped names.

If you specify a range of subinterfaces, you can specify a matching range of mapped names. Follow these guidelines for ranges:

- The mapped name must consist of an alphabetic portion followed by a numeric portion. The alphabetic portion of the mapped name must match for both ends of the range. For example, enter the following range:

int0-int10

If you enter **gigabitethernet0/1.1-gigabitethernet0/1.5 happy1-sad5**, for example, the command fails.

- The numeric portion of the mapped name must include the same quantity of numbers as the subinterface range. For example, both ranges include 100 interfaces:

gigabitethernet0/0.100-gigabitethernet0/0.199 int1-int100

If you enter **gigabitethernet0/0.100-gigabitethernet0/0.199 int1-int15**, for example, the command fails.

Specify **visible** to see physical interface properties in the **show interface** command even if you set a mapped name. The default **invisible** keyword specifies to only show the mapped name.

The following example shows gigabitethernet0/1.100, gigabitethernet0/1.200, and gigabitethernet0/2.300 through gigabitethernet0/1.305 assigned to the context. The mapped names are int1 through int8.

```
hostname(config-ctx)# allocate-interface gigabitethernet0/1.100 int1
hostname(config-ctx)# allocate-interface gigabitethernet0/1.200 int2
hostname(config-ctx)# allocate-interface gigabitethernet0/2.300-gigabitethernet0/2.305
int3-int8
```

- Step 4** To identify the URL from which the system downloads the context configuration, enter the following command:

```
hostname(config-ctx)# config-url url
```

When you add a context URL, the system immediately loads the context so that it is running, if the configuration is available.



Note

Enter the **allocate-interface** command(s) before you enter the **config-url** command. The security appliance must assign interfaces to the context before it loads the context configuration; the context configuration might include commands that refer to interfaces (**interface**, **nat**, **global**...). If you enter the **config-url** command first, the security appliance loads the context configuration immediately. If the context contains any commands that refer to interfaces, those commands fail.

See the following URL syntax:

- **disk:***/[path]/filename*

This URL indicates the internal Flash memory. The filename does not require a file extension, although we recommend using “.cfg”. If the configuration file is not available, you see the following message:

```
WARNING: Could not fetch the URL disk://url
INFO: Creating context with default config
```

You can then change to the context, configure it at the CLI, and enter the **write memory** command to write the file to Flash memory.



Note The admin context file must be stored on the internal Flash memory.

- **ftp:***//[user[:password]]@[server[:port]]/[path]/filename[;type=xx]*

The **type** can be one of the following keywords:

- **ap**—ASCII passive mode
- **an**—ASCII normal mode
- **ip**—(Default) Binary passive mode
- **in**—Binary normal mode

The server must be accessible from the admin context. The filename does not require a file extension, although we recommend using “.cfg”. If the configuration file is not available, you see the following message:

```
WARNING: Could not fetch the URL ftp://url
INFO: Creating context with default config
```

You can then change to the context, configure it at the CLI, and enter the **write memory** command to write the file to the FTP server.

- **http[s]:***//[user[:password]]@[server[:port]]/[path]/filename*

The server must be accessible from the admin context. The filename does not require a file extension, although we recommend using “.cfg”. If the configuration file is not available, you see the following message:

```
WARNING: Could not fetch the URL http://url
INFO: Creating context with default config
```

If you change to the context and configure the context at the CLI, you cannot save changes back to HTTP or HTTPS servers using the **write memory** command. You can, however, use the **copy tftp** command to copy the running configuration to a TFTP server.

- **tftp://[user[:password]@]server[:port]/[path/]filename[;int=interface_name]**

The server must be accessible from the admin context. Specify the interface name if you want to override the route to the server address. The filename does not require a file extension, although we recommend using “.cfg”. If the configuration file is not available, you see the following message:

```
WARNING: Could not fetch the URL tftp://url
INFO: Creating context with default config
```

You can then change to the context, configure it at the CLI, and enter the **write memory** command to write the file to the TFTP server.

To change the URL, reenter the **config-url** command with a new URL.

See the “[Changing the Security Context URL](#)” section on page 6-13 for more information about changing the URL.

For example, enter the following command:

```
hostname(config-ctx) # config-url tftp://joe:passw0rd1@10.1.1.1/configlets/test.cfg
```

- Step 5** (Optional) To assign the context to a resource class, enter the following command:

```
hostname(config-ctx) # member class_name
```

If you do not specify a class, the context belongs to the default class. You can only assign a context to one resource class.

For example, to assign the context to the gold class, enter the following command:

```
hostname(config-ctx) # member gold
```

- Step 6** To view context information, see the **show context** command in the *Cisco Security Appliance Command Reference*.

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

```
hostname(config) # admin-context administrator
hostname(config) # context administrator
hostname(config-ctx) # allocate-interface gigabitethernet0/0.1
hostname(config-ctx) # allocate-interface gigabitethernet0/1.1
hostname(config-ctx) # config-url flash:/admin.cfg

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

hostname(config-ctx) # context sample
hostname(config-ctx) # allocate-interface gigabitethernet0/1.200 int1
hostname(config-ctx) # allocate-interface gigabitethernet0/1.212 int2
hostname(config-ctx) # allocate-interface gigabitethernet0/1.230-gigabitethernet0/1.235
int3-int8
```

```
hostname(config-ctx) # config-url ftp://user1:passw0rd@10.1.1.1/configlets/sample.cfg
hostname(config-ctx) # member silver
```

Automatically Assigning MAC Addresses to Context Interfaces

To allow contexts to share interfaces, we suggest that you assign unique MAC addresses to each context interface. The MAC address is used to classify packets within a context. If you share an interface, but do not have unique MAC addresses for the interface in each context, then the destination IP address is used to classify packets. The destination address is matched with the context NAT configuration, and this method has some limitations compared to the MAC address method. See the [“How the Security Appliance Classifies Packets” section on page 3-3](#) for information about classifying packets.

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.

You can automatically assign private MAC addresses to each shared context interface by entering the following command in the system configuration:

```
hostname(config) # mac-address auto
```

For use with failover, the security appliance generates both an active and standby MAC address for each interface. If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption.

When you assign an interface to a context, the new MAC address is generated immediately. If you enable this command after you create context interfaces, then MAC addresses are generated for all interfaces immediately after you enter the command. If you use the **no mac-address auto** command, 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.

The MAC address is generated using the following format:

- Active unit MAC address: 12_slot.port_subid.contextid.
- Standby unit MAC address: 02_slot.port_subid.contextid.

For platforms with no interface slots, the slot is always 0. The *port* is the interface port. The *subid* is an internal ID for the subinterface, which is not viewable. The *contextid* is an internal ID for the context, viewable with the **show context detail** command. For example, the interface GigabitEthernet 0/1.200 in the context with the ID 1 has the following generated MAC addresses, where the internal ID for subinterface 200 is 31:

- Active: 1200.0131.0001
- Standby: 0200.0131.0001

In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, you can manually set the MAC address for the interface within the context. See the [“Configuring the Interface” section on page 7-2](#) to manually set the MAC address.

Changing Between Contexts and the System Execution Space

If you log in to the system execution space (or the admin context using Telnet or SSH), you can change between contexts and perform configuration and monitoring tasks within each context. The running configuration that you edit in a configuration mode, or that is used in the **copy** or **write** commands,

depends on your location. When you are in the system execution space, the running configuration consists only of the system configuration; when you are in a context, the running configuration consists only of that context. For example, you cannot view all running configurations (system plus all contexts) by entering the **show running-config** command. Only the current configuration displays.

To change between the system execution space and a context, or between contexts, see the following commands:

- To change to a context, enter the following command:

```
hostname# changeto context name
```

The prompt changes to the following:

```
hostname/name#
```

- To change to the system execution space, enter the following command:

```
hostname/admin# changeto system
```

The prompt changes to the following:

```
hostname#
```

Managing Security Contexts

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

- [Removing a Security Context, page 6-12](#)
- [Changing the Admin Context, page 6-13](#)
- [Changing the Security Context URL, page 6-13](#)
- [Reloading a Security Context, page 6-14](#)
- [Monitoring Security Contexts, page 6-15](#)

Removing a Security Context

You can only remove a context by editing the system configuration. You cannot remove the current admin context, unless you remove all contexts using the **clear context** command.



Note

If you use failover, there is a delay between when you remove the context on the active unit and when the context is removed on the standby unit. You might see an error message indicating that the number of interfaces on the active and standby units are not consistent; this error is temporary and can be ignored.

Use the following commands for removing contexts:

- To remove a single context, enter the following command in the system execution space:

```
hostname(config)# no context name
```

All context commands are also removed.

- To remove all contexts (including the admin context), enter the following command in the system execution space:


```
hostname(config)# clear context
```

Changing the Admin Context

The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses one of the contexts that is designated as the admin context.

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

You can set any context to be the admin context, as long as the configuration file is stored in the internal Flash memory. To set the admin context, enter the following command in the system execution space:

```
hostname(config)# admin-context context_name
```

Any remote management sessions, such as Telnet, SSH, or HTTPS, that are connected to the admin context are terminated. You must reconnect to the new admin context.



Note

A few system commands, including **ntp server**, identify an interface name that belongs to the admin context. If you change the admin context, and that interface name does not exist in the new admin context, be sure to update any system commands that refer to the interface.

Changing the Security Context URL

You cannot change the security context URL without reloading the configuration from the new URL.

The security appliance 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.

To change the URL for a context, perform the following steps:

-
- Step 1** If you do not want to merge the configuration, change to the context and clear its configuration by entering the following commands. If you want to perform a merge, skip to Step 2.

```
hostname# changeto context name
hostname/name# configure terminal
hostname/name(config)# clear configure all
```

- Step 2** If required, change to the system execution space by entering the following command:

```
hostname/name(config)# changeto system
```

- Step 3** To enter the context configuration mode for the context you want to change, enter the following command:

```
hostname(config)# context name
```

- Step 4** To enter the new URL, enter the following command:

```
hostname(config)# config-url new_url
```

The system immediately loads the context so that it is running.

Reloading a Security Context

You can reload the context in two ways:

- Clear the running configuration and then import the startup configuration.
This action clears most attributes associated with the context, such as connections and NAT tables.
- Remove the context from the system configuration.
This action clears additional attributes, such as memory allocation, which might be useful for troubleshooting. However, to add the context back to the system requires you to respecify the URL and interfaces.

This section includes the following topics:

- [Reloading by Clearing the Configuration, page 6-14](#)
- [Reloading by Removing and Re-adding the Context, page 6-15](#)

Reloading by Clearing the Configuration

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

- Step 1** To change to the context that you want to reload, enter the following command:

```
hostname# changeto context name
```

- Step 2** To access configuration mode, enter the following command:

```
hostname/name# configure terminal
```

- Step 3** To clear the running configuration, enter the following command:

```
hostname/name(config)# clear configure all
```

This command clears all connections.

- Step 4** To reload the configuration, enter the following command:

```
hostname/name(config)# copy startup-config running-config
```

The security appliance copies the configuration from the URL specified in the system configuration. You cannot change the URL from within a context.

Reloading by Removing and Re-adding the Context

To reload the context by removing the context and then re-adding it, perform the steps in the following sections:

1. [“Automatically Assigning MAC Addresses to Context Interfaces” section on page 6-11](#)
2. [“Configuring a Security Context” section on page 6-7](#)

Monitoring Security Contexts

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

- [Viewing Context Information, page 6-15](#)
- [Viewing Resource Allocation, page 6-16](#)
- [Viewing Resource Usage, page 6-19](#)
- [Monitoring SYN Attacks in Contexts, page 6-20](#)

Viewing Context Information

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

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

```
hostname# show context [name | detail | count]
```

The **detail** option shows additional information. See the following sample displays below for more information.

If you want to show information for a particular context, specify the *name*.

The **count** option shows the total number of contexts.

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

```
hostname# show context

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

[Table 6-2](#) shows each field description.

Table 6-2 *show context Fields*

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

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

```
hostname# show context detail

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

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

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

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

See the *Cisco Security Appliance Command Reference* for more information about the **detail** output.

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

```
hostname# show context count
Total active contexts: 2
```

Viewing Resource Allocation

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

To view the resource allocation, enter the following command:

```
hostname# show resource allocation [detail]
```

This command shows the resource allocation, but does not show the actual resources being used. See the [“Viewing Resource Usage” section on page 6-19](#) for more information about actual resource usage.

The **detail** argument shows additional information. See the following sample displays for more information.

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

```
hostname# show resource allocation

Resource      Total      % of Avail
-----
Conns [rate]  35000      N/A
Inspects [rate] 35000      N/A
Syslogs [rate] 10500      N/A
Conns          305000     30.50%
Hosts          78842      N/A
```

SSH	35	35.00%
Telnet	35	35.00%
Xlates	91749	N/A
All	unlimited	

Table 6-3 shows each field description.

Table 6-3 *show resource allocation Fields*

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

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

```
hostname# show resource allocation detail
Resource Origin:
  A Value was derived from the resource 'all'
  C Value set in the definition of this class
  D Value set in default class
```

Resource	Class	Mmbrs	Origin	Limit	Total	Total %
Conns [rate]	default	all	CA	unlimited		
	gold	1	C	34000	34000	N/A
	silver	1	CA	17000	17000	N/A
	bronze	0	CA	8500		
	All Contexts:	3			51000	N/A
Inspects [rate]	default	all	CA	unlimited		
	gold	1	DA	unlimited		
	silver	1	CA	10000	10000	N/A
	bronze	0	CA	5000		
	All Contexts:	3			10000	N/A
Syslogs [rate]	default	all	CA	unlimited		
	gold	1	C	6000	6000	N/A
	silver	1	CA	3000	3000	N/A
	bronze	0	CA	1500		
	All Contexts:	3			9000	N/A
Conns	default	all	CA	unlimited		
	gold	1	C	200000	200000	20.00%
	silver	1	CA	100000	100000	10.00%
	bronze	0	CA	50000		
	All Contexts:	3			300000	30.00%
Hosts	default	all	CA	unlimited		
	gold	1	DA	unlimited		
	silver	1	CA	26214	26214	N/A
	bronze	0	CA	13107		
	All Contexts:	3			26214	N/A
SSH	default	all	C	5		
	gold	1	D	5	5	5.00%

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

Table 6-4 shows each field description.

Table 6-4 *show resource allocation detail Fields*

Field	Description
Resource	The name of the resource that you can limit.
Class	The name of each class, including the default class. The All contexts field shows the total values across all classes.
Mmbrs	The number of contexts assigned to each class.
Origin	The origin of the resource limit, as follows: - A—You set this limit with the all option, instead of as an individual resource. - C—This limit is derived from the member class. - D—This limit was not defined in the member class, but was derived from the default class. For a context assigned to the default class, the value will be “C” instead of “D.” The security appliance can combine “A” with “C” or “D.”
Limit	The limit of the resource per context, as an absolute number. If you specified a percentage in the class definition, the security appliance converts the percentage to an absolute number for this display.
Total	The total amount of the resource that is allocated across all contexts in the class. The amount is an absolute number of concurrent instances or instances per second. If the resource is unlimited, this display is blank.
% of Avail	The percentage of the total system resources that is allocated across all contexts in the class. If the resource is unlimited, this display is blank. If the resource does not have a system limit, then this column shows N/A.

Viewing Resource Usage

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

From the system execution space, view the resource usage for each context by entering the following command:

```
hostname# show resource usage [context context_name | top n | all | summary | system]
[resource {resource_name | all} | detail] [counter counter_name [count_threshold]]
```

By default, **all** context usage is displayed; each context is listed separately.

Enter the **top n** keyword to show the contexts that are the top *n* users of the specified resource. You must specify a single resource type, and not **resource all**, with this option.

The **summary** option shows all context usage combined.

The **system** option shows all context usage combined, but shows the system limits for resources instead of the combined context limits.

For the **resource resource_name**, see [Table 6-1](#) for available resource names. See also the **show resource type** command. Specify **all** (the default) for all types.

The **detail** option shows the resource usage of all resources, including those you cannot manage. For example, you can view the number of TCP intercepts.

The **counter counter_name** is one of the following keywords:

- **current**—Shows the active concurrent instances or the current rate of the resource.
- **denied**—Shows the number of instances that were denied because they exceeded the resource limit shown in the Limit column.
- **peak**—Shows the peak concurrent instances, or the peak rate of the resource since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **all**—(Default) Shows all statistics.

The *count_threshold* sets the number above which resources are shown. The default is 1. If the usage of the resource is below the number you set, then the resource is not shown. If you specify **all** for the counter name, then the *count_threshold* applies to the current usage.



Note

To show all resources, set the *count_threshold* to 0.

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

```
hostname# show resource usage context admin
```

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

The following is sample output from the **show resource usage summary** command, which shows the resource usage for all contexts and all resources. This sample shows the limits for 6 contexts.

```
hostname# show resource usage summary
```

Resource	Current	Peak	Limit	Denied	Context
Syslogs [rate]	1743	2132	N/A	0	Summary
Conns	584	763	280000 (S)	0	Summary

```

Xlates                8526          8966          N/A              0 Summary
Hosts                 254           254           N/A              0 Summary
Conns [rate]          270           535           N/A             1704 Summary
Inspects [rate]        270           535           N/A              0 Summary
S = System: Combined context limits exceed the system limit; the system limit is shown.

```

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

```
hostname# show resource usage summary
```

```

Resource            Current      Peak      Limit    Denied    Context
Telnet               1          1      100[S]      0      Summary
SSH                  2          2      100[S]      0      Summary
Conns                56         90       N/A         0      Summary
Hosts                89        102       N/A         0      Summary
S = System: Combined context limits exceed the system limit; the system limit is shown.

```

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

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

```

Resource            Current      Peak      Limit    Denied    Context
Telnet              0          0       100         0      System
SSH                 0          0       100         0      System
ASDM                0          0        32         0      System
Syslogs [rate]      1         18       N/A         0      System
Conns               0          1    280000         0      System
Xlates              0          0       N/A         0      System
Hosts               0          2       N/A         0      System
Conns [rate]        1          1       N/A         0      System
Inspects [rate]     0          0       N/A         0      System

```

Monitoring SYN Attacks in Contexts

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

You can monitor the rate of attacks for individual contexts using the **show perfmon** command; you can monitor the amount of resources being used by TCP intercept for individual contexts using the **show resource usage detail** command; you can monitor the resources being used by TCP intercept for the entire system using the **show resource usage summary detail** command.

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

```
hostname/admin# show perfmon
```

```

Context:admin
PERFMON STATS:   Current      Average
Xlates           0/s          0/s

```


Connections	0/s	0/s
TCP Conns	0/s	0/s
UDP Conns	0/s	0/s
URL Access	0/s	0/s
URL Server Req	0/s	0/s
WebSns Req	0/s	0/s
TCP Fixup	0/s	0/s
HTTP Fixup	0/s	0/s
FTP Fixup	0/s	0/s
AAA Authen	0/s	0/s
AAA Author	0/s	0/s
AAA Account	0/s	0/s
TCP Intercept	322779/s	322779/s

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

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

Resource	Current	Peak	Limit	Denied	Context
memory	843732	847288	unlimited	0	admin
chunk:channels	14	15	unlimited	0	admin
chunk:fixup	15	15	unlimited	0	admin
chunk:hole	1	1	unlimited	0	admin
chunk:ip-users	10	10	unlimited	0	admin
chunk:list-elem	21	21	unlimited	0	admin
chunk:list-hdr	3	4	unlimited	0	admin
chunk:route	2	2	unlimited	0	admin
chunk:static	1	1	unlimited	0	admin
tcp-intercepts	328787	803610	unlimited	0	admin
np-statics	3	3	unlimited	0	admin
statics	1	1	unlimited	0	admin
ace-rules	1	1	unlimited	0	admin
console-access-rul	2	2	unlimited	0	admin
fixup-rules	14	15	unlimited	0	admin
memory	959872	960000	unlimited	0	c1
chunk:channels	15	16	unlimited	0	c1
chunk:dbgtrace	1	1	unlimited	0	c1
chunk:fixup	15	15	unlimited	0	c1
chunk:global	1	1	unlimited	0	c1
chunk:hole	2	2	unlimited	0	c1
chunk:ip-users	10	10	unlimited	0	c1
chunk:udp-ctrl-blk	1	1	unlimited	0	c1
chunk:list-elem	24	24	unlimited	0	c1
chunk:list-hdr	5	6	unlimited	0	c1
chunk:nat	1	1	unlimited	0	c1
chunk:route	2	2	unlimited	0	c1
chunk:static	1	1	unlimited	0	c1
tcp-intercept-rate	16056	16254	unlimited	0	c1
globals	1	1	unlimited	0	c1
np-statics	3	3	unlimited	0	c1
statics	1	1	unlimited	0	c1
nats	1	1	unlimited	0	c1
ace-rules	2	2	unlimited	0	c1
console-access-rul	2	2	unlimited	0	c1
fixup-rules	14	15	unlimited	0	c1
memory	232695716	232020648	unlimited	0	system
chunk:channels	17	20	unlimited	0	system
chunk:dbgtrace	3	3	unlimited	0	system
chunk:fixup	15	15	unlimited	0	system
chunk:ip-users	4	4	unlimited	0	system
chunk:list-elem	1014	1014	unlimited	0	system
chunk:list-hdr	1	1	unlimited	0	system
chunk:route	1	1	unlimited	0	system

```

block:16384          510          885 unlimited          0 system
block:2048           32           34 unlimited          0 system

```

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

```

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

```



CHAPTER 7

Configuring Interface Parameters

This chapter describes how to configure each interface and subinterface for a name, security level, and IP address. For single context mode, the procedures in this chapter continue the interface configuration started in [Chapter 5, “Configuring Ethernet Settings and Subinterfaces.”](#) For multiple context mode, the procedures in [Chapter 5, “Configuring Ethernet Settings and Subinterfaces,”](#) are performed in the system execution space, while the procedures in this chapter are performed within each security context.



Note

To configure interfaces for the ASA 5505 adaptive security appliance, see [Chapter 4, “Configuring Switch Ports and VLAN Interfaces for the Cisco ASA 5505 Adaptive Security Appliance.”](#)

This chapter includes the following sections:

- [Security Level Overview, page 7-1](#)
- [Configuring the Interface, page 7-2](#)
- [Allowing Communication Between Interfaces on the Same Security Level, page 7-6](#)

Security Level Overview

Each interface must have a security level from 0 (lowest) to 100 (highest). For example, you should assign your most secure network, such as the inside host network, to level 100. While the outside network connected to the Internet can be level 0. Other networks, such as DMZs can be in between. You can assign interfaces to the same security level. See the [“Allowing Communication Between Interfaces on the Same Security Level”](#) section on page 7-6 for more information.

The level controls the following behavior:

- Network access—By default, there is an implicit permit from a higher security interface to a lower security interface (outbound). Hosts on the higher security interface can access any host on a lower security interface. You can limit access by applying an access list to the interface.

If you enable communication for same security interfaces (see the [“Allowing Communication Between Interfaces on the Same Security Level”](#) section on page 7-6), there is an implicit permit for interfaces to access other interfaces on the same security level or lower.

- Inspection engines—Some application inspection engines are dependent on the security level. For same security interfaces, inspection engines apply to traffic in either direction.
 - NetBIOS inspection engine—Applied only for outbound connections.

- SQL*Net inspection engine—If a control connection for the SQL*Net (formerly OraServ) port exists between a pair of hosts, then only an inbound data connection is permitted through the security appliance.
- Filtering—HTTP(S) and FTP filtering applies only for outbound connections (from a higher level to a lower level).

For some security interfaces, you can filter traffic in either direction.

- NAT control—When you enable NAT control, you must configure NAT for hosts on a higher security interface (inside) when they access hosts on a lower security interface (outside).

Without NAT control, or for some security interfaces, you can choose to use NAT between any interface, or you can choose not to use NAT. Keep in mind that configuring NAT for an outside interface might require a special keyword.

- **established** command—This command allows return connections from a lower security host to a higher security host if there is already an established connection from the higher level host to the lower level host.

For some security interfaces, you can configure **established** commands for both directions.

Configuring the Interface

By default, all physical interfaces are shut down. You must enable the physical interface before any traffic can pass through an enabled subinterface. For multiple context mode, if you allocate a physical interface or subinterface to a context, the interfaces are enabled by default in the context. However, before traffic can pass through the context interface, you must also enable the interface in the system configuration. If you shut down an interface in the system execution space, then that interface is down in all contexts that share it.

Before you can complete your configuration and allow traffic through the security appliance, you need to configure an interface name, and for routed mode, an IP address. You should also change the security level from the default, which is 0. If you name an interface “inside” and you do not set the security level explicitly, then the security appliance sets the security level to 100.



Note

If you are using failover, do not use this procedure to name interfaces that you are reserving for failover and Stateful Failover communications. See [Chapter 14, “Configuring Failover.”](#) to configure the failover and state links.

For multiple context mode, follow these guidelines:

- Configure the context interfaces from within each context.
- You can only configure context interfaces that you already assigned to the context in the system configuration.
- The system configuration only lets you configure Ethernet settings and VLANs. The exception is for failover interfaces; do not configure failover interfaces with this procedure. See the Failover chapter for more information.



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.

To configure an interface or subinterface, perform the following steps:

Step 1 To specify the interface you want to configure, enter the following command:

```
hostname(config)# interface {physical_interface[.subinterface] | mapped_name}
```

The *physical_interface* ID includes the type, slot, and port number as *type[slot/port]*.

The physical interface types include the following:

- **ethernet**
- **gigabitethernet**

For the PIX 500 series security appliance, enter the type followed by the port number, for example, **ethernet0**.

For the ASA 5500 series adaptive security appliance, enter the type followed by slot/port, for example, **gigabitethernet0/1**. Interfaces that are built into the chassis are assigned to slot 0, while interfaces on the 4GE SSM are assigned to slot 1. For the ASA 5550 adaptive security appliance, for maximum throughput, be sure to balance your traffic over the two interface slots; for example, assign the inside interface to slot 1 and the outside interface to slot 0.

The ASA 5510 and higher adaptive security appliance also includes the following type:

- **management**

The management interface is a Fast Ethernet interface designed for management traffic only, and is specified as **management0/0**. You can, however, use it for through traffic if desired (see the **management-only** command). In transparent firewall mode, you can use the management interface in addition to the two interfaces allowed for through traffic. You can also add subinterfaces to the management interface to provide management in each security context for multiple context mode.

Append the *subinterface* ID to the physical interface ID separated by a period (.).

In multiple context mode, enter the mapped name if one was assigned using the **allocate-interface** command.

For example, enter the following command:

```
hostname(config)# interface gigabitethernet0/1.1
```

Step 2 To name the interface, enter the following command:

```
hostname(config-if)# nameif 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. Do not enter the **no** form, because that command causes all commands that refer to that name to be deleted.

Step 3 To set the security level, enter the following command:

```
hostname(config-if)# security-level number
```

Where *number* is an integer between 0 (lowest) and 100 (highest).

Step 4 (Optional) To set an interface to management-only mode, enter the following command:

```
hostname(config-if)# management-only
```

The ASA 5510 and higher adaptive security appliance includes a dedicated management interface called Management 0/0, which is meant to support traffic to the security appliance. However, you can configure any interface to be a management-only interface using the **management-only** command. Also, for Management 0/0, you can disable management-only mode so the interface can pass through traffic just like any other interface.

**Note**

Transparent firewall mode allows only two interfaces to pass through traffic; however, on the ASA 5510 and higher adaptive security appliance, you can use the Management 0/0 interface (either the physical interface or a subinterface) as a third interface for management traffic. The mode is not configurable in this case and must always be management-only.

Step 5 To set the IP address, enter one of the following commands.

In routed firewall mode, you set the IP address for all interfaces. In transparent firewall mode, you do not set the IP address for each interface, but rather for the whole security appliance or context. The exception is for the Management 0/0 management-only interface, which does not pass through traffic. To set the management IP address for transparent firewall mode, see the [“Setting the Management IP Address for a Transparent Firewall” section on page 8-5](#). To set the IP address of the Management 0/0 interface or subinterface, use one of the following commands.

To set an IPv6 address, see the [“Configuring IPv6 on an Interface” section on page 12-3](#).

For failover, you must set the IP address an standby address manually; DHCP and PPPoE are not supported.

- To set the IP address manually, enter the following command:

```
hostname(config-if)# ip address ip_address [mask] [standby ip_address]
```

The **standby** keyword and address is used for failover. See [Chapter 14, “Configuring Failover,”](#) for more information.

- To obtain an IP address from a DHCP server, enter the following command:

```
hostname(config-if)# ip address dhcp [setroute]
```

Reenter this command to reset the DHCP lease and request a new lease.

If you do not enable the interface using the **no shutdown** command before you enter the **ip address dhcp** command, some DHCP requests might not be sent.

- To obtain an IP address from a PPPoE server, see [Chapter 35, “Configuring the PPPoE Client.”](#)

Step 6 (Optional) To assign a private MAC address to this interface, enter the following command:

```
hostname(config-if)# mac-address mac_address [standby mac_address]
```

The *mac_address* is in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 00C.F142.4CDE.

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 use with failover, set the **standby** MAC address. If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.

In multiple context mode, if you share an interface between contexts, you can assign a unique MAC address to the interface in each context. This feature lets the security appliance easily classify packets into the appropriate context. Using a shared interface without unique MAC addresses is possible, but has some limitations. See the [“How the Security Appliance Classifies Packets” section on page 3-3](#) for more information. You can assign each MAC address manually, or you can automatically generate MAC addresses for shared interfaces in contexts. See the [“Automatically Assigning MAC Addresses to Context Interfaces” section on page 6-11](#) to automatically generate MAC addresses. If you automatically generate MAC addresses, you can use the **mac-address** command 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.

Step 7 To enable the interface, if it is not already enabled, enter the following command:

```
hostname(config-if)# no shutdown
```

To disable the interface, enter the **shutdown** command. If you enter the **shutdown** command for a physical interface, you also shut down all subinterfaces. If you shut down an interface in the system execution space, then that interface is shut down in all contexts that share it, even though the context configurations show the interface as enabled.

The following example configures parameters for the physical interface in single mode:

```
hostname(config)# interface gigabitethernet0/1
hostname(config-if)# speed 1000
hostname(config-if)# duplex full
hostname(config-if)# nameif inside
hostname(config-if)# security-level 100
hostname(config-if)# ip address 10.1.1.1 255.255.255.0
hostname(config-if)# no shutdown
```

The following example configures parameters for a subinterface in single mode:

```
hostname(config)# interface gigabitethernet0/1.1
hostname(config-subif)# vlan 101
hostname(config-subif)# nameif dmz1
hostname(config-subif)# security-level 50
hostname(config-subif)# ip address 10.1.2.1 255.255.255.0
hostname(config-subif)# mac-address 000C.F142.4CDE standby 020C.F142.4CDE
hostname(config-subif)# no shutdown
```

The following example configures interface parameters in multiple context mode for the system configuration, and allocates the gigabitethernet 0/1.1 subinterface to contextA:

```
hostname(config)# interface gigabitethernet0/1
hostname(config-if)# speed 1000
hostname(config-if)# duplex full
hostname(config-if)# no shutdown
hostname(config-if)# interface gigabitethernet0/1.1
hostname(config-subif)# vlan 101
hostname(config-subif)# no shutdown
hostname(config-subif)# context contextA
hostname(config-ctx)# ...
hostname(config-ctx)# allocate-interface gigabitethernet0/1.1
```

The following example configures parameters in multiple context mode for the context configuration:

```
hostname/contextA(config)# interface gigabitethernet0/1.1
hostname/contextA(config-if)# nameif inside
hostname/contextA(config-if)# security-level 100
hostname/contextA(config-if)# ip address 10.1.2.1 255.255.255.0
hostname/contextA(config-if)# mac-address 030C.F142.4CDE standby 040C.F142.4CDE
hostname/contextA(config-if)# no shutdown
```

Allowing Communication Between Interfaces on the Same Security Level

By default, interfaces on the same security level cannot communicate with each other. Allowing communication between same security interfaces provides the following benefits:

- You can configure more than 101 communicating interfaces.
If you use different levels for each interface and do not assign any interfaces to the same security level, you can configure only one interface per level (0 to 100).
- You want traffic to flow freely between all same security interfaces without access lists.

**Note**

If you enable NAT control, you do not need to configure NAT between same security level interfaces. See the [“NAT and Same Security Level Interfaces” section on page 17-13](#) for more information on NAT and same security level interfaces.

If you enable same security interface communication, you can still configure interfaces at different security levels as usual.

To enable interfaces on the same security level so that they can communicate with each other, enter the following command:

```
hostname(config)# same-security-traffic permit inter-interface
```

To disable this setting, use the **no** form of this command.



CHAPTER 8

Configuring Basic Settings

This chapter describes how to configure basic settings on your security appliance that are typically required for a functioning configuration. This chapter includes the following sections:

- [Changing the Login Password, page 8-1](#)
- [Changing the Enable Password, page 8-1](#)
- [Setting the Hostname, page 8-2](#)
- [Setting the Domain Name, page 8-2](#)
- [Setting the Date and Time, page 8-2](#)
- [Setting the Management IP Address for a Transparent Firewall, page 8-5](#)

Changing the Login Password

The login password is used for Telnet and SSH connections. By default, the login password is “cisco.” To change the password, enter the following command:

```
hostname(config)# {passwd | password} password
```

You can enter **passwd** or **password**. The password is a case-sensitive password of up to 16 alphanumeric and special characters. You can use any character in the password except a question mark or a space.

The password is saved in the configuration in encrypted form, so you cannot view the original password after you enter it. Use the `no password` command to restore the password to the default setting.

Changing the Enable Password

The enable password lets you enter privileged EXEC mode. By default, the enable password is blank. To change the enable password, enter the following command:

```
hostname(config)# enable password password
```

The *password* is a case-sensitive password of up to 16 alphanumeric and special characters. You can use any character in the password except a question mark or a space.

This command changes the password for the highest privilege level. If you configure local command authorization, you can set enable passwords for each privilege level from 0 to 15.

The password is saved in the configuration in encrypted form, so you cannot view the original password after you enter it. Enter the **enable password** command without a password to set the password to the default, which is blank.

Setting the Hostname

When you set a hostname for the security appliance, that name appears in the command line prompt. If you establish sessions to multiple devices, the hostname helps you keep track of where you enter commands. The default hostname depends on your platform.

For multiple context mode, the hostname that you set in the system execution space appears in the command line prompt for all contexts. The hostname that you optionally set within a context does not appear in the command line, but can be used by the **banner** command **\$(hostname)** token.

To specify the hostname for the security appliance or for a context, enter the following command:

```
hostname(config)# hostname name
```

This name can be up to 63 characters. A hostname must start and end with a letter or digit, and have as interior characters only letters, digits, or a hyphen.

This name appears in the command line prompt. For example:

```
hostname(config)# hostname farscape  
farscape(config)#
```

Setting the Domain Name

The security appliance appends the domain name as a suffix to unqualified names. For example, if you set the domain name to “example.com,” and specify a syslog server by the unqualified name of “jupiter,” then the security appliance qualifies the name to “jupiter.example.com.”

The default domain name is default.domain.invalid.

For multiple context mode, you can set the domain name for each context, as well as within the system execution space.

To specify the domain name for the security appliance, enter the following command:

```
hostname(config)# domain-name name
```

For example, to set the domain as example.com, enter the following command:

```
hostname(config)# domain-name example.com
```

Setting the Date and Time

This section describes how to set the date and time, either manually or dynamically using an NTP server. Time derived from an NTP server overrides any time set manually. This section also describes how to set the time zone and daylight saving time date range.

**Note**

In multiple context mode, set the time in the system configuration only.

This section includes the following topics:

- [Setting the Time Zone and Daylight Saving Time Date Range, page 8-3](#)
- [Setting the Date and Time Using an NTP Server, page 8-4](#)
- [Setting the Date and Time Manually, page 8-5](#)

Setting the Time Zone and Daylight Saving Time Date Range

By default, the time zone is UTC and the daylight saving time date range is from 2:00 a.m. on the first Sunday in April to 2:00 a.m. on the last Sunday in October. To change the time zone and daylight saving time date range, perform the following steps:

- Step 1** To set the time zone, enter the following command in global configuration mode:

```
hostname(config)# clock timezone zone [-]hours [minutes]
```

Where *zone* specifies the time zone as a string, for example, **PST** for Pacific Standard Time.

The *[-]hours* value sets the number of hours of offset from UTC. For example, PST is **-8** hours.

The *minutes* value sets the number of minutes of offset from UTC.

- Step 2** To change the date range for daylight saving time from the default, enter one of the following commands.

The default recurring date range is from 2:00 a.m. on the first Sunday in April to 2:00 a.m. on the last Sunday in October.

- To set the start and end dates for daylight saving time as a specific date in a specific year, enter the following command:

```
hostname(config)# clock summer-time zone date {day month | month day} year hh:mm {day month | month day} year hh:mm [offset]
```

If you use this command, you need to reset the dates every year.

The *zone* value specifies the time zone as a string, for example, **PDT** for Pacific Daylight Time.

The *day* value sets the day of the month, from 1 to 31. You can enter the day and month as **April 1** or as **1 April**, for example, depending on your standard date format.

The *month* value sets the month as a string. You can enter the day and month as **April 1** or as **1 April**, for example, depending on your standard date format.

The *year* value sets the year using four digits, for example, **2004**. The year range is 1993 to 2035.

The *hh:mm* value sets the hour and minutes in 24-hour time.

The *offset* value sets the number of minutes to change the time for daylight saving time. By default, the value is 60 minutes.

- To specify the start and end dates for daylight saving time, in the form of a day and time of the month, and not a specific date in a year, enter the following command.

```
hostname(config)# clock summer-time zone recurring [week weekday month hh:mm week weekday month hh:mm] [offset]
```

This command lets you set a recurring date range that you do not need to alter yearly.

The *zone* value specifies the time zone as a string, for example, **PDT** for Pacific Daylight Time.

The *week* value specifies the week of the month as an integer between 1 and 4 or as the words **first** or **last**. For example, if the day might fall in the partial fifth week, then specify **last**.

The *weekday* value specifies the day of the week: **Monday, Tuesday, Wednesday**, and so on.

The *month* value sets the month as a string.

The *hh:mm* value sets the hour and minutes in 24-hour time.

The *offset* value sets the number of minutes to change the time for daylight saving time. By default, the value is 60 minutes.

Setting the Date and Time Using an NTP Server

To obtain the date and time from an NTP server, perform the following steps:

Step 1 To configure authentication with an NTP server, perform the following steps:

- a. To enable authentication, enter the following command:

```
hostname(config)# ntp authenticate
```

- b. To specify an authentication key ID to be a trusted key, which is required for authentication with an NTP server, enter the following command:

```
hostname(config)# ntp trusted-key key_id
```

Where the *key_id* is between 1 and 4294967295. You can enter multiple trusted keys for use with multiple servers.

- c. To set a key to authenticate with an NTP server, enter the following command:

```
hostname(config)# ntp authentication-key key_id md5 key
```

Where *key_id* is the ID you set in Step 1b using the **ntp trusted-key** command, and *key* is a string up to 32 characters in length.

Step 2 To identify an NTP server, enter the following command:

```
hostname(config)# ntp server ip_address [key key_id] [source interface_name] [prefer]
```

Where the *key_id* is the ID you set in [Step 1b](#) using the **ntp trusted-key** command.

The **source interface_name** identifies the outgoing interface for NTP packets if you do not want to use the default interface in the routing table. Because the system does not include any interfaces in multiple context mode, specify an interface name defined in the admin context.

The **prefer** keyword sets this NTP server as the preferred server if multiple servers have similar accuracy. NTP uses an algorithm to determine which server is the most accurate and synchronizes to that one. If servers are of similar accuracy, then the **prefer** keyword specifies which of those servers to use. However, if a server is significantly more accurate than the preferred one, the security appliance uses the more accurate one. For example, the security appliance uses a server of stratum 2 over a server of stratum 3 that is preferred.

You can identify multiple servers; the security appliance uses the most accurate server.



Note

SNTP is not supported; only NTP is supported.

Setting the Date and Time Manually

To set the date time manually, enter the following command:

```
hostname# clock set hh:mm:ss {month day | day month} year
```

Where *hh:mm:ss* sets the hour, minutes, and seconds in 24-hour time. For example, set **20:54:00** for 8:54 pm.

The *day* value sets the day of the month, from 1 to 31. You can enter the day and month as **april 1** or as **1 april**, for example, depending on your standard date format.

The *month* value sets the month. Depending on your standard date format, you can enter the day and month as **april 1** or as **1 april**.

The *year* value sets the year using four digits, for example, **2004**. The year range is 1993 to 2035.

The default time zone is UTC. If you change the time zone after you enter the **clock set** command using the **clock timezone** command, the time automatically adjusts to the new time zone.

This command sets the time in the hardware chip, and does not save the time in the configuration file. This time endures reboots. Unlike the other **clock** commands, this command is a privileged EXEC command. To reset the clock, you need to set a new time for the **clock set** command.

Setting the Management IP Address for a Transparent Firewall

Transparent firewall mode only

A transparent firewall does not participate in IP routing. The only IP configuration required for the security appliance is to set the management IP address. This address is required because the security appliance uses this address as the source address for traffic originating on the security appliance, such as system messages or communications with AAA servers. You can also use this address for remote management access.

For multiple context mode, set the management IP address within each context.

To set the management IP address, enter the following command:

```
hostname(config)# ip address ip_address [mask] [standby ip_address]
```

This address must be on the same subnet as the upstream and downstream routers. You cannot set the subnet to a host subnet (255.255.255.255). This address must be IPv4; the transparent firewall does not support IPv6.

The **standby** keyword and address is used for failover. See [Chapter 14, “Configuring Failover,”](#) for more information.



CHAPTER 9

Configuring IP Routing

This chapter describes how to configure IP routing on the security appliance. This chapter includes the following sections:

- [How Routing Behaves Within the ASA Security Appliance, page 9-1](#)
- [Configuring Static and Default Routes, page 9-2](#)
- [Defining Route Maps, page 9-7](#)
- [Configuring OSPF, page 9-8](#)
- [Configuring RIP, page 9-20](#)
- [The Routing Table, page 9-24](#)
- [Dynamic Routing and Failover, page 9-26](#)

How Routing Behaves Within the ASA Security Appliance

The ASA security appliance uses both routing table and XLATE tables for routing decisions. To handle destination IP translated traffic, that is, untranslated traffic, ASA searches for existing XLATE, or static translation to select the egress interface. The selection process is as follows:

Egress Interface Selection Process

1. If 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 destination IP translating XLATE does not exist, but a matching static translation exists, then the egress interface is determined from the static route and an XLATE is created, and the routing table is not used.
3. If destination IP translating XLATE does not exist and no matching static translation exists, the packet is not destination IP translated. The security appliance processes this packet by looking up the route to select 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 egress interface using any method described above, an additional route lookup is performed to find out suitable next hop(s) that belong to previously selected egress interface. If there are no routes in routing table that explicitly belong to selected interface, the packet is dropped with level 6 error message 110001 "no route to host", even if there is another route for a given destination network that belongs to different egress interface. If the route that belongs to selected egress interface is found, the packet is forwarded to corresponding next hop.

Load sharing on the security appliance is possible only for multiple next-hops available using single egress interface. Load sharing cannot share multiple egress interfaces.

If dynamic routing is in use on security appliance and route table changes after XLATE creation, for example route flap, then destination translated traffic is still forwarded using old XLATE, not via route table, until XLATE times out. It may be either forwarded to wrong interface or dropped with message 110001 "no route to host" if old route was removed from the old interface and attached to another one by routing process.

The same problem may happen when there is no route flaps on the security appliance itself, but some routing process is flapping around it, sending source translated packets that belong to the same flow through the security appliance using different interfaces. Destination translated return packets may be forwarded back using the wrong egress interface.

This issue has a high probability in same security traffic configuration, where virtually any traffic may be either source-translated or destination-translated, depending on direction of 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. XLATE timeout may be decreased if necessary. To ensure that this rarely happens, make sure that there is no route flaps on security appliance and around it. That is, ensure that destination translated packets that belong to the same flow are always forwarded the same way through the security appliance.

Configuring Static and Default Routes

This section describes how to configure static and default routes on the security appliance.

Multiple context mode does not support dynamic routing, so you must use static routes for any networks to which the security appliance is not directly connected; for example, when there is a router between a network and the security appliance.

You might want to use static routes in single context mode in the following cases:

- Your networks use a different router discovery protocol from RIP or OSPF.
- Your network is small and you can easily manage static routes.
- You do not want the traffic or CPU overhead associated with routing protocols.

The simplest option is to configure a default route to send all traffic to an upstream router, relying on the router to route the traffic for you. However, in some cases the default gateway might not be able to reach the destination network, so you must also configure more specific static routes. For example, if the default gateway is outside, then the default route cannot direct traffic to any inside networks that are not directly connected to the security appliance.

In transparent firewall mode, for traffic that originates on the security appliance and is destined for a non-directly connected network, you need to configure either a default route or static routes so the security appliance knows out of which interface to send traffic. Traffic that originates on the security

appliance 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.

The security appliance supports up to three equal cost routes on the same interface for load balancing.

This section includes the following topics:

- [Configuring a Static Route, page 9-3](#)
- [Configuring a Default Route, page 9-4](#)
- [Configuring Static Route Tracking, page 9-5](#)

For information about configuring IPv6 static and default routes, see the “[Configuring IPv6 Default and Static Routes](#)” section on page 12-5.

Configuring a Static Route

To add a static route, enter the following command:

```
hostname(config)# route if_name dest_ip mask gateway_ip [distance]
```

The *dest_ip* and *mask* is the IP address for the destination network and the *gateway_ip* is the address of the next-hop router. The addresses you specify for the static route are the addresses that are in the packet before entering the security appliance and performing NAT.

The *distance* is the administrative distance for the route. The default is 1 if you do not specify a value. Administrative distance is a parameter used to compare routes among different routing protocols. The default administrative distance for static routes is 1, giving it precedence over routes discovered by dynamic routing protocols but not directly 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.

Static routes remain in the routing table even if the specified gateway becomes unavailable. If the specified gateway becomes unavailable, you need to remove the static route from the routing table manually. However, static routes are removed from the routing table if the specified interface goes down. They are reinstated when the interface comes back up.



Note

If you create a static route with an administrative distance greater than the administrative distance of the routing protocol running on the security appliance, then a route to the specified destination discovered by the routing protocol takes precedence over the static route. The static route is used only if the dynamically discovered route is removed from the routing table.

The following example creates a static route that sends all traffic destined for 10.1.1.0/24 to the router (10.1.2.45) connected to the inside interface:

```
hostname(config)# route inside 10.1.1.0 255.255.255.0 10.1.2.45 1
```

You can define up to three equal cost routes to the same destination per interface. ECMP is not supported across multiple interfaces. With ECMP, the traffic is not necessarily divided evenly between the routes; traffic is distributed among the specified gateways based on an algorithm that hashes the source and destination IP addresses.

The following example shows static routes that are equal cost routes that direct traffic to three different gateways on the outside interface. The security appliance distributes the traffic among the specified gateways.

```
hostname(config)# route outside 10.10.10.0 255.255.255.0 192.168.1.1
hostname(config)# route outside 10.10.10.0 255.255.255.0 192.168.1.2
hostname(config)# route outside 10.10.10.0 255.255.255.0 192.168.1.3
```

Configuring a Default Route

A default route identifies the gateway IP address to which the security appliance sends all IP packets for which it does not have a learned or static route. A default route is simply a static route with 0.0.0.0/0 as the destination IP address. Routes that identify a specific destination take precedence over the default route.



Note

In ASA software Versions 7.0 and later, if you have two default routes configured on different interfaces that have different metrics, the connection to the ASA firewall that is made from the higher metric interface fails, but connections to the ASA firewall from the lower metric interface succeed as expected. PIX software Version 6.3 supports connections from both the the higher and the lower metric interfaces.

You can define up to three equal cost default route entries per device. Defining more than one equal cost default route entry causes the traffic sent to the default route to be distributed among the specified gateways. When defining more than one default route, you must specify the same interface for each entry.

If you attempt to define more than three equal cost default routes, or if you attempt to define a default route with a different interface than a previously defined default route, you receive the message “ERROR: Cannot add route entry, possible conflict with existing routes.”

You can define a separate default route for tunneled traffic along with the standard default route. When you create a default route with the **tunneled** option, all traffic from a tunnel terminating on the security appliance that cannot be routed using learned or static routes, is sent to this route. For traffic emerging from a tunnel, this route overrides over any other configured or learned default routes.

The following restrictions apply to default routes with the **tunneled** option:

- Do not enable unicast RPF (**ip verify reverse-path**) on the egress interface of tunneled route. Enabling uRPF on the egress interface of a tunneled route causes the session to fail.
- Do not enable TCP intercept on the egress interface of the tunneled route. Doing so causes the session to fail.
- Do not use the VoIP inspection engines (CTIQBE, H.323, GTP, MGCP, RTSP, SIP, SKINNY), the DNS inspect engine, or the DCE RPC inspection engine with tunneled routes. These inspection engines ignore the tunneled route.

You cannot define more than one default route with the tunneled option; ECMP for tunneled traffic is not supported.

To define the default route, enter the following command:

```
hostname(config)# route if_name 0.0.0.0 0.0.0.0 gateway_ip [distance | tunneled]
```



Tip

You can enter 0 0 instead of 0.0.0.0 0.0.0.0 for the destination network address and mask, for example:

```
hostname(config)# route outside 0 0 192.168.1 1
```

The following example shows a security appliance configured with three equal cost default routes and a default route for tunneled traffic. Unencrypted traffic received by the security appliance for which there is no static or learned route is distributed among the gateways with the IP addresses 192.168.2.1, 192.168.2.2, 192.168.2.3. Encrypted traffic received by the security appliance for which there is no static or learned route is passed to the gateway with the IP address 192.168.2.4.

```
hostname(config)# route outside 0 0 192.168.2.1
hostname(config)# route outside 0 0 192.168.2.2
hostname(config)# route outside 0 0 192.168.2.3
hostname(config)# route outside 0 0 192.168.2.4 tunneled
```

Configuring 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 security appliance goes down.

The static route tracking feature provides a method for tracking the availability of a static route and installing a backup route if the primary route should fail. This allows you to, for example, define a default route to an ISP gateway and a backup default route to a secondary ISP in case the primary ISP becomes unavailable.

The security appliance does this by associating a static route with a monitoring target that you define. It monitors the target using ICMP echo requests. If an echo reply is not received within a specified time period, the object is considered down and the associated route is removed from the routing table. A previously configured backup route is used in place of the removed route.

When selecting a monitoring target, you need to make sure it can respond to ICMP echo requests. The target can be any network object that you choose, but you should consider using:

- 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 security appliance needs to communicate with
- a persistent network object on the destination network (a desktop or notebook computer that may be shut down at night is not a good choice)

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 interface with route tracking.

To configure static route tracking, perform the following steps:

Step 1 Configure the tracked object monitoring parameters:

- a. Define the monitoring process:

```
hostname(config)# sla monitor sla_id
```

If you are configuring a new monitoring process, you are taken to SLA monitor configuration mode. If you are changing the monitoring parameters for an unscheduled monitoring process that already has a type defined, you are taken directly to the SLA protocol configuration mode.

- b. Specify the monitoring protocol. If you are changing the monitoring parameters for an unscheduled monitoring process that already has a type defined, you are taken directly to SLA protocol configuration mode and cannot change this setting.

```
hostname(config-sla-monitor)# type echo protocol ipIcmpEcho target_ip interface
if_name
```

The *target_ip* is the IP address of the network object whose availability the tracking process monitors. While this object is available, the tracking process route is installed in the routing table. When this object becomes unavailable, the tracking process removed the route and the backup route is used in its place.

c. Schedule the monitoring process:

```
hostname(config)# sla monitor schedule sla_id [life {forever | seconds}] [start-time
{hh:mm[:ss] [month day | day month] | pending | now | after hh:mm:ss}] [ageout
seconds] [recurring]
```

Typically, you will use **sla monitor schedule sla_id life forever start-time now** for the monitoring schedule, and allow the monitoring configuration determine how often the testing occurs. However, you can schedule this monitoring process to begin in the future and to only occur at specified times.

Step 2 Associate a tracked static route with the SLA monitoring process by entering the following command:

```
hostname(config)# track track_id rtr sla_id reachability
```

The *track_id* is a tracking number you assign with this command. The *sla_id* is the ID number of the SLA process you defined in [Step 1](#).

Step 3 Define the static route to be installed in the routing table while the tracked object is reachable using one of the following options:

- To track a static route, enter the following command:

```
hostname(config)# route if_name dest_ip mask gateway_ip [admin_distance] track
track_id
```

You cannot use the **tunneled** option with the **route** command with static route tracking.

- To track a default route obtained through DHCP, enter the following commands:

```
hostname(config)# interface phy_if
hostname(config-if)# dhcp client route track track_id
hostname(config-if)# ip addresss dhcp setroute
hostname(config-if)# exit
```



Note You must use the **setroute** argument with the **ip address dhcp** command to obtain the default route using DHCP.

- To track a default route obtained through PPPoE, enter the following commands:

```
hostname(config)# interface phy_if
hostname(config-if)# pppoe client route track track_id
hostname(config-if)# ip addresss pppoe setroute
hostname(config-if)# exit
```



Note You must use the **setroute** argument with the **ip address pppoe** command to obtain the default route using PPPoE.

Step 4 Define the backup route to use when the tracked object is unavailable using one of the following options. The administrative distance of the backup route must be greater than the administrative distance of the tracked route. If it is not, the backup route will be installed in the routing table instead of the tracked route.

- To use a static route, enter the following command:

```
hostname(config)# route if_name dest_ip mask gateway_ip [admin_distance]
```

The static route must have the same destination and mask as the tracked route. If you are tracking a default route obtained through DHCP or PPPoE, then the address and mask would be 0.0.0.0 0.0.0.0.

- To use a default route obtained through DHCP, enter the following commands:

```
hostname(config)# interface phy_if
hostname(config-if)# dhcp client route track track_id
hostname(config-if)# dhcp client route distance admin_distance
hostname(config-if)# ip addresss dhcp setroute
hostname(config-if)# exit
```

You must use the **setroute** argument with the **ip address dhcp** command to obtain the default route using DHCP. Make sure the administrative distance is greater than the administrative distance of the tracked route.

- To use a default route obtained through PPPoE, enter the following commands:

```
hostname(config)# interface phy_if
hostname(config-if)# pppoe client route track track_id
hostname(config-if)# pppoe client route distance admin_distance
hostname(config-if)# ip addresss pppoe setroute
hostname(config-if)# exit
```

You must use the **setroute** argument with the **ip address pppoe** command to obtain the default route using PPPoE. Make sure the administrative distance is greater than the administrative distance of the tracked route.

Defining Route Maps

Route maps are used when redistributing routes into an OSPF or RIP 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.

To define a route map, perform the following steps:

- Step 1** To create a route map entry, enter the following command:

```
hostname(config)# route-map name {permit | deny} [sequence_number]
```

Route map entries are read in order. You can identify the order using the *sequence_number* option, or the security appliance uses the order in which you add the entries.

- Step 2** Enter one or more **match** commands:

- To match any routes that have a destination network that matches a standard ACL, enter the following command:

```
hostname(config-route-map)# match ip address acl_id [acl_id] [...]
```

If you specify more than one ACL, then the route can match any of the ACLs.

- To match any routes that have a specified metric, enter the following command:

```
hostname(config-route-map)# match metric metric_value
```

The *metric_value* can be from 0 to 4294967295.

- To match any routes that have a next hop router address that matches a standard ACL, enter the following command:

```
hostname(config-route-map)# match ip next-hop acl_id [acl_id] [...]
```

If you specify more than one ACL, then the route can match any of the ACLs.

- To match any routes with the specified next hop interface, enter the following command:

```
hostname(config-route-map)# match interface if_name
```

If you specify more than one interface, then the route can match either interface.

- To match any routes that have been advertised by routers that match a standard ACL, enter the following command:

```
hostname(config-route-map)# match ip route-source acl_id [acl_id] [...]
```

If you specify more than one ACL, then the route can match any of the ACLs.

- To match the route type, enter the following command:

```
hostname(config-route-map)# match route-type {internal | external [type-1 | type-2]}
```

Step 3 Enter one or more **set** commands.

If a route matches the **match** commands, then the following **set** commands determine the action to perform on the route before redistributing it.

- To set the metric, enter the following command:

```
hostname(config-route-map)# set metric metric_value
```

The *metric_value* can be a value between 0 and 294967295

- To set the metric type, enter the following command:

```
hostname(config-route-map)# set metric-type {type-1 | type-2}
```

The following example shows how to redistribute routes with a hop count equal to 1 into OSPF. The security appliance redistributes these routes as external LSAs with a metric of 5, metric type of Type 1.

```
hostname(config)# route-map 1-to-2 permit
hostname(config-route-map)# match metric 1
hostname(config-route-map)# set metric 5
hostname(config-route-map)# set metric-type type-1
```

Configuring OSPF

This section describes how to configure OSPF. This section includes the following topics:

- [OSPF Overview, page 9-9](#)
- [Enabling OSPF, page 9-10](#)
- [Redistributing Routes Into OSPF, page 9-10](#)
- [Configuring OSPF Interface Parameters, page 9-11](#)
- [Configuring OSPF Area Parameters, page 9-13](#)

- [Configuring OSPF NSSA, page 9-14](#)
- [Defining Static OSPF Neighbors, page 9-16](#)
- [Configuring Route Summarization Between OSPF Areas, page 9-15](#)
- [Configuring Route Summarization When Redistributing Routes into OSPF, page 9-16](#)
- [Generating a Default Route, page 9-17](#)
- [Configuring Route Calculation Timers, page 9-17](#)
- [Logging Neighbors Going Up or Down, page 9-18](#)
- [Displaying OSPF Update Packet Pacing, page 9-19](#)
- [Monitoring OSPF, page 9-19](#)
- [Restarting the OSPF Process, page 9-20](#)

OSPF Overview

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 security appliance 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 security appliance 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 security appliance supports the following OSPF features:

- Support of intra-area, interarea, and external (Type I and Type II) routes.
- Support of a virtual link.
- OSPF LSA flooding.
- Authentication to OSPF packets (both password and MD5 authentication).
- Support for configuring the security appliance as a designated router or a designated backup router. The security appliance also can be set up as an ABR; however, the ability to configure the security appliance as an ASBR is limited to default information only (for example, injecting a default route).
- Support for stub areas and not-so-stubby-areas.

- Area boundary router type-3 LSA filtering.
- Advertisement of static and global address translations.

Enabling OSPF

To enable OSPF, you need to create an OSPF routing process, specify the range of IP addresses associated with the routing process, then assign area IDs associated with that range of IP addresses.

To enable OSPF, perform the following steps:

-
- Step 1** To create an OSPF routing process, enter the following command:

```
hostname(config)# router ospf process_id
```

This command enters the router configuration mode for this OSPF process.

The *process_id* is an internally used identifier for this routing process. It can be any positive integer. This ID does not have to match the ID on any other device; it is for internal use only. You can use a maximum of two processes.

- Step 2** To define the IP addresses on which OSPF runs and to define the area ID for that interface, enter the following command:

```
hostname(config-router)# network ip_address mask area area_id
```

The following example shows how to enable OSPF:

```
hostname(config)# router ospf 2
hostname(config-router)# network 10.0.0.0 255.0.0.0 area 0
```

Redistributing Routes Into OSPF

The security appliance can control the redistribution of routes between OSPF routing processes. The security appliance matches and changes routes according to settings in the **redistribute** command or by using a route map. See also the [“Generating a Default Route” section on page 9-17](#) for another use for route maps.

To redistribute static, connected, RIP, or OSPF routes into an OSPF process, perform the following steps:

-
- Step 1** (Optional) Create a route-map to further define which routes from the specified routing protocol are redistributed into the OSPF routing process. See the [“Defining Route Maps” section on page 9-7](#).
- Step 2** If you have not already done so, enter the router configuration mode for the OSPF process you want to redistribute into by entering the following command:

```
hostname(config)# router ospf process_id
```

- Step 3** To specify the routes you want to redistribute, enter the following command:

```
hostname(config-router)# redistribute {ospf process_id
[match {internal | external 1 | external 2}] | static | connected | rip}
[metric metric-value] [metric-type {type-1 | type-2}] [tag tag_value] [subnets] [route-map
map_name]
```


The **ospf process_id**, **static**, **connected**, and **rip** keywords specify from where you want to redistribute routes.

You can either use the options in this command to match and set route properties, or you can use a route map. The **tag** and **subnets** options do not have equivalents in the **route-map** command. If you use both a route map and options in the **redistribute** command, then they must match.

The following example shows route redistribution from OSPF process 1 into OSPF process 2 by matching routes with a metric equal to 1. The security appliance redistributes these routes as external LSAs with a metric of 5, metric type of Type 1, and a tag equal to 1.

```
hostname(config)# route-map 1-to-2 permit
hostname(config-route-map)# match metric 1
hostname(config-route-map)# set metric 5
hostname(config-route-map)# set metric-type type-1
hostname(config-route-map)# set tag 1
hostname(config-route-map)# router ospf 2
hostname(config-router)# redistribute ospf 1 route-map 1-to-2
```

The following example shows the specified OSPF process routes being redistributed into OSPF process 109. The OSPF metric is remapped to 100.

```
hostname(config)# router ospf 109
hostname(config-router)# redistribute ospf 108 metric 100 subnets
```

The following example shows route redistribution where the link-state cost is specified as 5 and the metric type is set to external, indicating that it has lower priority than internal metrics.

```
hostname(config)# router ospf 1
hostname(config-router)# redistribute ospf 2 metric 5 metric-type external
```

Configuring OSPF Interface Parameters

You can alter some interface-specific OSPF parameters as necessary. You are not required to alter any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: **ospf hello-interval**, **ospf dead-interval**, and **ospf authentication-key**. Be sure that if you configure any of these parameters, the configurations for all routers on your network have compatible values.

To configure OSPF interface parameters, perform the following steps:

Step 1 To enter the interface configuration mode, enter the following command:

```
hostname(config)# interface interface_name
```

Step 2 Enter any of the following commands:

- To specify the authentication type for an interface, enter the following command:

```
hostname(config-interface)# ospf authentication [message-digest | null]
```

- To assign a password to be used by neighboring OSPF routers on a network segment that is using the OSPF simple password authentication, enter the following command:

```
hostname(config-interface)# ospf authentication-key key
```

The *key* can be any continuous string of characters up to 8 bytes in length.

The password created by this command is used as a key that is inserted directly into the OSPF header when the security appliance software originates routing protocol packets. A separate password can be assigned to each network on a per-interface basis. All neighboring routers on the same network must have the same password to be able to exchange OSPF information.

- To explicitly specify the cost of sending a packet on an OSPF interface, enter the following command:

```
hostname(config-interface)# ospf cost cost
```

The *cost* is an integer from 1 to 65535.

- To set the number of seconds that a device must wait before it declares a neighbor OSPF router down because it has not received a hello packet, enter the following command:

```
hostname(config-interface)# ospf dead-interval seconds
```

The value must be the same for all nodes on the network.

- To specify the length of time between the hello packets that the security appliance sends on an OSPF interface, enter the following command:

```
hostname(config-interface)# ospf hello-interval seconds
```

The value must be the same for all nodes on the network.

- To enable OSPF MD5 authentication, enter the following command:

```
hostname(config-interface)# ospf message-digest-key key_id md5 key
```

Set the following values:

- *key_id*—An identifier in the range from 1 to 255.
- *key*—Alphanumeric password of up to 16 bytes.

Usually, one key per interface is used to generate authentication information when sending packets and to authenticate incoming packets. The same key identifier on the neighbor router must have the same key value.

We recommend that you not keep more than one key per interface. Every time you add a new key, you should remove the old key to prevent the local system from continuing to communicate with a hostile system that knows the old key. Removing the old key also reduces overhead during rollover.

- To set the priority to help determine the OSPF designated router for a network, enter the following command:

```
hostname(config-interface)# ospf priority number_value
```

The *number_value* is between 0 to 255.

- To specify the number of seconds between LSA retransmissions for adjacencies belonging to an OSPF interface, enter the following command:

```
hostname(config-interface)# ospf retransmit-interval seconds
```

The *seconds* must be greater than the expected round-trip delay between any two routers on the attached network. The range is from 1 to 65535 seconds. The default is 5 seconds.

- To set the estimated number of seconds required to send a link-state update packet on an OSPF interface, enter the following command:

```
hostname(config-interface)# ospf transmit-delay seconds
```

The *seconds* is from 1 to 65535 seconds. The default is 1 second.

The following example shows how to configure the OSPF interfaces:

```
hostname(config)# router ospf 2
hostname(config-router)# network 2.0.0.0 255.0.0.0 area 0
hostname(config-router)# interface inside
hostname(config-interface)# ospf cost 20
hostname(config-interface)# ospf retransmit-interval 15
hostname(config-interface)# ospf transmit-delay 10
hostname(config-interface)# ospf priority 20
hostname(config-interface)# ospf hello-interval 10
hostname(config-interface)# ospf dead-interval 40
hostname(config-interface)# ospf authentication-key cisco
hostname(config-interface)# ospf message-digest-key 1 md5 cisco
hostname(config-interface)# ospf authentication message-digest
```

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

```
hostname(config)# show ospf

Routing Process "ospf 2" with ID 20.1.89.2 and Domain ID 0.0.0.2
Supports only single TOS(TOS0) routes
Supports opaque LSA
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
Number of external LSA 5. Checksum Sum 0x 26da6
Number of opaque AS LSA 0. Checksum Sum 0x      0
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
External flood list length 0
  Area BACKBONE(0)
    Number of interfaces in this area is 1
    Area has no authentication
    SPF algorithm executed 2 times
    Area ranges are
    Number of LSA 5. Checksum Sum 0x 209a3
    Number of opaque link LSA 0. Checksum Sum 0x      0
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
```

Configuring OSPF Area Parameters

You can configure several area parameters. These area parameters (shown in the following task table) include setting authentication, defining stub areas, and assigning specific costs to the default summary route. Authentication provides password-based protection against unauthorized access to an area.

Stub areas are areas into which information on external routes is not sent. Instead, there is a default external route generated by the ABR, into the stub area for destinations outside the autonomous system. To take advantage of the OSPF stub area support, default routing must be used in the stub area. To further reduce the number of LSAs sent into a stub area, you can configure the **no-summary** keyword of the **area stub** command on the ABR to prevent it from sending summary link advertisement (LSA type 3) into the stub area.

To specify area parameters for your network, perform the following steps:

- Step 1** If you have not already done so, enter the router configuration mode for the OSPF process you want to configure by entering the following command:

```
hostname(config)# router ospf process_id
```

- Step 2** Enter any of the following commands:

- To enable authentication for an OSPF area, enter the following command:

```
hostname(config-router)# area area-id authentication
```

- To enable MD5 authentication for an OSPF area, enter the following command:

```
hostname(config-router)# area area-id authentication message-digest
```

- To define an area to be a stub area, enter the following command:

```
hostname(config-router)# area area-id stub [no-summary]
```

- To assign a specific cost to the default summary route used for the stub area, enter the following command:

```
hostname(config-router)# area area-id default-cost cost
```

The *cost* is an integer from 1 to 65535. The default is 1.

The following example shows how to configure the OSPF area parameters:

```
hostname(config)# router ospf 2
hostname(config-router)# area 0 authentication
hostname(config-router)# area 0 authentication message-digest
hostname(config-router)# area 17 stub
hostname(config-router)# area 17 default-cost 20
```

Configuring OSPF NSSA

The OSPF implementation of an NSSA is similar to an OSPF 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 OSPF to a remote site that is using a different routing protocol using NSSA.

Before the implementation of NSSA, the connection between the corporate site border router and the remote router could not be run as an OSPF 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 OSPF to cover the remote connection by defining the area between the corporate router and the remote router as an NSSA.

To specify area parameters for your network as needed to configure OSPF NSSA, perform the following steps:

- Step 1** If you have not already done so, enter the router configuration mode for the OSPF process you want to configure by entering the following command:

```
hostname(config)# router ospf process_id
```

- Step 2** Enter any of the following commands:

- To define an NSSA area, enter the following command:

```
hostname(config-router)# area area-id nssa [no-redistribution]  
[default-information-originate]
```

- To summarize groups of addresses, enter the following command:

```
hostname(config-router)# summary address ip_address mask [not-advertise] [tag tag]
```

This command helps reduce the size of the routing table. Using this command for OSPF causes an OSPF ASBR to advertise one external route as an aggregate for all redistributed routes that are covered by the address.

OSPF does not support **summary-address 0.0.0.0 0.0.0.0**.

In the following example, the summary address 10.1.0.0 includes address 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. Only the address 10.1.0.0 is advertised in an external link-state advertisement:

```
hostname(config-router)# summary-address 10.1.1.0 255.255.0.0
```

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 will not be able to communicate.

Configuring Route Summarization Between OSPF 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 such that they are contiguous, you can configure the area boundary router to advertise a summary route that covers all the individual networks within the area that fall into the specified range.

To define an address range for route summarization, perform the following steps:

- Step 1** If you have not already done so, enter the router configuration mode for the OSPF process you want to configure by entering the following command:

```
hostname(config)# router ospf process_id
```

Step 2 To set the address range, enter the following command:

```
hostname(config-router)# area area-id range ip-address mask [advertise | not-advertise]
```

The following example shows how to configure route summarization between OSPF areas:

```
hostname(config)# router ospf 1  
hostname(config-router)# area 17 range 12.1.0.0 255.255.0.0
```

Configuring Route Summarization When Redistributing Routes into OSPF

When routes from other protocols are redistributed into OSPF, each route is advertised individually in an external LSA. However, you can configure the security appliance to advertise a single route for all the redistributed routes that are covered by a specified network address and mask. This configuration decreases the size of the OSPF link-state database.

To configure the software advertisement on one summary route for all redistributed routes covered by a network address and mask, perform the following steps:

Step 1 If you have not already done so, enter the router configuration mode for the OSPF process you want to configure by entering the following command:

```
hostname(config)# router ospf process_id
```

Step 2 To set the summary address, enter the following command:

```
hostname(config-router)# summary-address ip_address mask [not-advertise] [tag tag]
```



Note OSPF does not support **summary-address 0.0.0.0 0.0.0.0**.

The following example shows how to configure route summarization. The summary address 10.1.0.0 includes address 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. Only the address 10.1.0.0 is advertised in an external link-state advertisement:

```
hostname(config)# router ospf 1  
hostname(config-router)# summary-address 10.1.0.0 255.255.0.0
```

Defining Static OSPF Neighbors

You need to define static OSPF neighbors to advertise OSPF routes over a point-to-point, non-broadcast network. This lets you broadcast OSPF advertisements across an existing VPN connection without having to encapsulate the advertisements in a GRE tunnel.

To define a static OSPF neighbor, perform the following tasks:

Step 1 Create a static route to the OSPF neighbor. See the [“Configuring Static and Default Routes”](#) section on [page 9-2](#) for more information about creating static routes.

Step 2 Define the OSPF neighbor by performing the following tasks:

- a. Enter router configuration mode for the OSPF process. Enter the following command:

```
hostname(config)# router ospf pid
```

- b. Define the OSPF neighbor by entering the following command:

```
hostname(config-router)# neighbor addr [interface if_name]
```

The *addr* argument is the IP address of the OSPF neighbor. The *if_name* is the interface used to communicate with the neighbor. If the OSPF neighbor is not on the same network as any of the directly-connected interfaces, you must specify the **interface**.

Generating a Default Route

You can force an autonomous system boundary router to generate a default route into an OSPF routing domain. Whenever you specifically configure redistribution of routes into an OSPF routing domain, the router automatically becomes an autonomous system boundary router. However, an autonomous system boundary router does not by default generate a default route into the OSPF routing domain.

To generate a default route, perform the following steps:

Step 1 If you have not already done so, enter the router configuration mode for the OSPF process you want to configure by entering the following command:

```
hostname(config)# router ospf process_id
```

Step 2 To force the autonomous system boundary router to generate a default route, enter the following command:

```
hostname(config-router)# default-information originate [always] [metric metric-value]  
[metric-type {1 | 2}] [route-map map-name]
```

The following example shows how to generate a default route:

```
hostname(config)# router ospf 2  
hostname(config-router)# default-information originate always
```

Configuring Route Calculation Timers

You can configure the delay time between when OSPF receives a topology change and when it starts an SPF calculation. You also can configure the hold time between two consecutive SPF calculations.

To configure route calculation timers, perform the following steps:

Step 1 If you have not already done so, enter the router configuration mode for the OSPF process you want to configure by entering the following command:

```
hostname(config)# router ospf process_id
```

- Step 2** To configure the route calculation time, enter the following command:

```
hostname(config-router)# timers spf spf-delay spf-holdtime
```

The *spf-delay* is the delay time (in seconds) between when OSPF receives a topology change and when it starts an SPF calculation. It can be an integer from 0 to 65535. The default time is 5 seconds. A value of 0 means that there is no delay; that is, the SPF calculation is started immediately.

The *spf-holdtime* is the minimum time (in seconds) between two consecutive SPF calculations. It can be an integer from 0 to 65535. The default time is 10 seconds. A value of 0 means that there is no delay; that is, two SPF calculations can be done, one immediately after the other.

The following example shows how to configure route calculation timers:

```
hostname(config)# router ospf 1  
hostname(config-router)# timers spf 10 120
```

Logging Neighbors Going Up or Down

By default, the system sends a system message when an OSPF neighbor goes up or down.

Configure this command if you want to know about OSPF neighbors going up or down without turning on the **debug ospf adjacency** command. The **log-adj-changes** router configuration command provides a higher level view of the peer relationship with less output. Configure **log-adj-changes detail** if you want to see messages for each state change.

To log neighbors going up or down, perform the following steps:

-
- Step 1** If you have not already done so, enter the router configuration mode for the OSPF process you want to configure by entering the following command:

```
hostname(config)# router ospf process_id
```

- Step 2** To configure logging for neighbors going up or down, enter the following command:

```
hostname(config-router)# log-adj-changes [detail]
```



Note Logging must be enabled for the the neighbor up/down messages to be sent.

The following example shows how to log neighbors up/down messages:

```
hostname(config)# router ospf 1  
hostname(config-router)# log-adj-changes detail
```


Displaying OSPF Update Packet Pacing

OSPF update packets are automatically paced so they are not sent less than 33 milliseconds apart. Without pacing, some update packets could get lost in situations where the link is slow, a neighbor could not receive the updates quickly enough, or the router could run out of buffer space. For example, without pacing packets might be dropped if either of the following topologies exist:

- A fast router is connected to a slower router over a point-to-point link.
- During flooding, several neighbors send updates to a single router at the same time.

Pacing is also used between resends to increase efficiency and minimize lost retransmissions. You also can display the LSAs waiting to be sent out an interface. The benefit of the pacing is that OSPF update and retransmission packets are sent more efficiently.

There are no configuration tasks for this feature; it occurs automatically.

To observe OSPF packet pacing by displaying a list of LSAs waiting to be flooded over a specified interface, enter the following command:

```
hostname# show ospf flood-list if_name
```

Monitoring OSPF

You can display specific statistics such as the contents of IP routing tables, caches, and databases. You can 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 display various OSPF routing statistics, perform one of the following tasks, as needed:

- To display general information about OSPF routing processes, enter the following command:

```
hostname# show ospf [process-id [area-id]]
```
- To display the internal OSPF routing table entries to the ABR and ASBR, enter the following command:

```
hostname# show ospf border-routers
```
- To display lists of information related to the OSPF database for a specific router, enter the following command:

```
hostname# show ospf [process-id [area-id]] database
```
- To display a list of LSAs waiting to be flooded over an interface (to observe OSPF packet pacing), enter the following command:

```
hostname# show ospf flood-list if-name
```
- To display OSPF-related interface information, enter the following command:

```
hostname# show ospf interface [if_name]
```
- To display OSPF neighbor information on a per-interface basis, enter the following command:

```
hostname# show ospf neighbor [interface-name] [neighbor-id] [detail]
```
- To display a list of all LSAs requested by a router, enter the following command:

```
hostname# show ospf request-list neighbor if_name
```

- To display a list of all LSAs waiting to be resent, enter the following command:

```
hostname# show ospf retransmission-list neighbor if_name
```

- To display a list of all summary address redistribution information configured under an OSPF process, enter the following command:

```
hostname# show ospf [process-id] summary-address
```

- To display OSPF-related virtual links information, enter the following command:

```
hostname# show ospf [process-id] virtual-links
```

Restarting the OSPF Process

To restart an OSPF process, clear redistribution, or counters, enter the following command:

```
hostname(config)# clear ospf pid {process | redistribution | counters  
[neighbor [neighbor-interface] [neighbor-id]]}
```

Configuring RIP

Devices that support RIP send routing-update messages at regular intervals and when the network topology changes. These RIP packets contain information about the networks that the devices can reach, as well as the number of routers or gateways that a packet must travel through to reach the destination address. RIP generates more traffic than OSPF, but is easier to configure.

RIP has advantages over static routes because the initial configuration is simple, and you do not need to update the configuration when the topology changes. The disadvantage to RIP is that there is more network and processing overhead than static routing.

The security appliance supports RIP Version 1 and RIP Version 2.

This section describes how to configure RIP. This section includes the following topics:

- [Enabling and Configuring RIP, page 9-20](#)
- [Redistributing Routes into the RIP Routing Process, page 9-22](#)
- [Configuring RIP Send/Receive Version on an Interface, page 9-22](#)
- [Enabling RIP Authentication, page 9-23](#)
- [Monitoring RIP, page 9-23](#)

Enabling and Configuring RIP

You can only enable one RIP routing process on the security appliance. After you enable the RIP routing process, you must define the interfaces that will participate in that routing process using the **network** command. By default, the security appliance sends RIP Version 1 updates and accepts RIP Version 1 and Version 2 updates.

To enable and configure the RIP routing process, perform the following steps:

-
- Step 1** Start the RIP routing process by entering the following command in global configuration mode:
- ```
hostname(config): router rip
```
- You enter router configuration mode for the RIP routing process.
- Step 2** Specify the interfaces that will participate in the RIP routing process. Enter the following command for each interface that will participate in the RIP routing process:
- ```
hostname(config-router): network network_address
```
- If an interface belongs to a network defined by this command, the interface will participate in the RIP routing process. If an interface does not belong to a network defined by this command, it will not send or receive RIP updates.
- Step 3** (Optional) Specify the version of RIP used by the security appliance by entering the following command:
- ```
hostname(config-router): version [1 | 2]
```
- You can override this setting on a per-interface basis.
- Step 4** (Optional) To generate a default route into RIP, enter the following command:
- ```
hostname(config-router): default-information originate
```
- Step 5** (Optional) To specify an interface to operate in passive mode, enter the following command:
- ```
hostname(config-router): passive-interface [default | if_name]
```
- Using the **default** keyword causes all interfaces to operate in passive mode. Specifying an interface name sets only that interface to passive RIP mode. In passive mode, RIP routing updates are accepted by but not sent out of the specified interface. You can enter this command for each interface you want to set to passive mode.
- Step 6** (Optional) Disable automatic route summarization by entering the following command:
- ```
hostname(config-router): no auto-summarize
```
- RIP Version 1 always uses automatic route summarization; you cannot disable it for RIP Version 1. RIP Version 2 uses route summarization by default; you can disable it using this command.
- Step 7** (Optional) To filter the networks received in updates, perform the following steps:
- Create a standard access list permitting the networks you want the RIP process to allow in the routing table and denying the networks you want the RIP process to discard.
 - Enter the following command to apply the filter. You can specify an interface to apply the filter to only those updates received by that interface.
- ```
hostname(config-router): distribute-list acl in [interface if_name]
```
- You can enter this command for each interface you want to apply a filter to. If you do not specify an interface name, the filter is applied to all RIP updates.
- Step 8** (Optional) To filter the networks sent in updates, perform the following steps:
- Create a standard access list permitting the networks you want the RIP process to advertise and denying the networks you do not want the RIP process to advertise.
  - Enter the following command to apply the filter. You can specify an interface to apply the filter to only those updates sent by that interface.
- ```
hostname(config-router): distribute-list acl out [interface if_name]
```

You can enter this command for each interface you want to apply a filter to. If you do not specify an interface name, the filter is applied to all RIP updates.

Redistributing Routes into the RIP Routing Process

You can redistribute routes from the OSPF, static, and connected routing processes into the RIP routing process.

To redistribute a routes into the RIP routing process, perform the following steps:

-
- Step 1** (Optional) Create a route-map to further define which routes from the specified routing protocol are redistributed in to the RIP routing process. See the [“Defining Route Maps” section on page 9-7](#) for more information about creating a route map.
- Step 2** Choose one of the following options to redistribute the selected route type into the RIP routing process.
- To redistribute connected routes into the RIP routing process, enter the following command:

```
hostname(config-router): redistribute connected [metric {metric_value | transparent}] [route-map map_name]
```
 - To redistribute static routes into the RIP routing process, enter the following command:

```
hostname(config-router): redistribute static [metric {metric_value | transparent}] [route-map map_name]
```
 - To redistribute routes from an OSPF routing process into the RIP routing process, enter the following command:

```
hostname(config-router): redistribute ospf pid [match {internal | external [1 | 2] | nssa-external [1 | 2]] [metric {metric_value | transparent}] [route-map map_name]
```
-

Configuring RIP Send/Receive Version on an Interface

You can override the globally-set version of RIP the security appliance uses to send and receive RIP updates on a per-interface basis.

To configure the RIP send and receive

-
- Step 1** (Optional) To specify the version of RIP advertisements sent from an interface, perform the following steps:
- Enter interface configuration mode for the interface you are configuring by entering the following command:

```
hostname(config)# interface phy_if
```
 - Specify the version of RIP to use when sending RIP updates out of the interface by entering the following command:

```
hostname(config-if)# rip send version {[1] [2]}
```

- Step 2** (Optional) To specify the version of RIP advertisements permitted to be received by an interface, perform the following steps:
- Enter interface configuration mode for the interface you are configuring by entering the following command:

```
hostname(config)# interface phy_if
```
 - Specify the version of RIP to allow when receiving RIP updates on the interface by entering the following command:

```
hostname(config-if)# rip receive version {[1] [2]}
```
- RIP updates received on the interface that do not match the allowed version are dropped.
-

Enabling RIP Authentication

The security appliance supports RIP message authentication for RIP Version 2 messages.

To enable RIP message authentication, perform the following steps:

-
- Step 1** Enter interface configuration mode for the interface you are configuring by entering the following command:

```
hostname(config)# interface phy_if
```
- Step 2** (Optional) Set the authentication mode by entering the following command. By default, text authentication is used. MD5 authentication is recommended.

```
hostname(config-if)# rip authentication mode {text | md5}
```
- Step 3** Enable authentication and configure the authentication key by entering the following command:

```
hostname(config-if)# rip authentication key key key_id key-id
```
-

Monitoring RIP

To display various RIP routing statistics, perform one of the following tasks, as needed:

- To display the contents of the RIP routing database, enter the following command:

```
hostname# show rip database
```
- To display the RIP commands in the running configuration, enter the following command:

```
hostname# show running-config router rip
```

Use the following **debug** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco TAC. Debugging output is assigned high priority in the CPU process and can render the system unusable. It is best to use **debug** commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that increased **debug** command processing overhead will affect system performance.

- To display RIP processing events, enter the following command:

```
hostname# debug rip events
```

- To display RIP database events, enter the following command:

```
hostname# debug rip database
```

The Routing Table

This section contains the following topics:

- [Displaying the Routing Table, page 9-24](#)
- [How the Routing Table is Populated, page 9-24](#)
- [How Forwarding Decisions are Made, page 9-26](#)

Displaying the Routing Table

To view the entries in the routing table, enter the following command:

```
hostname# show route
```

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
        i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
        * - candidate default, U - per-user static route, o - ODR
        P - periodic downloaded static route
```

```
Gateway of last resort is 10.86.194.1 to network 0.0.0.0
```

```
S    10.1.1.0 255.255.255.0 [3/0] via 10.86.194.1, outside
C    10.86.194.0 255.255.254.0 is directly connected, outside
S*   0.0.0.0 0.0.0.0 [1/0] via 10.86.194.1, outside
```

On the ASA 5505 adaptive security appliance, the following route is also shown. It is the internal loopback interface, which is used by the VPN Hardware Client feature for individual user authentication.

```
C 127.1.0.0 255.255.0.0 is directly connected, _internal_loopback
```

How the Routing Table is Populated

The security appliance routing table can be populated by statically defined routes, directly connected routes, and routes discovered by the RIP and OSPF routing protocols. Because the security appliance 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 in to 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 determine which route to use.

- If the security appliance 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 security appliance 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 is entered into the routing table.

Administrative distance is a route parameter that security appliance 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 9-1](#) shows the default administrative distance values for the routing protocols supported by the security appliance.

Table 9-1 Default Administrative Distance for Supported Routing Protocols

Route Source	Default Administrative Distance
Connected interface	0
Static route	1
OSPF	110
RIP	120

The smaller the administrative distance value, the more preference is given to the protocol. For example, if the security appliance 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 - 100), the security appliance chooses the OSPF route because OSPF has a higher preference. This means the router adds the OSPF version of the route to the routing table.

In the above example, if the source of the OSPF-derived route was lost (for example, due to a power shutdown), the security appliance 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 security appliance the command was entered on. The administrative distance is not advertised in routing updates.

Administrative distance does not affect the routing process. The OSPF and RIP routing processes only advertise the routes that have been discovered by the routing process or redistributed into the routing process. For example, the RIP routing process advertises RIP routes, even if routes discovered by the OSPF routing process are used in the security appliance 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 security appliance. 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 packets for that destination are distributed among the interfaces associated with that route.
- 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 a security appliance with the following routes in the routing table:

```
hostname# show route
....
R   192.168.32.0/24 [120/4] via 10.1.1.2
O   192.168.32.0/19 [110/229840] via 10.1.1.3
....
```

In this case, a packet destined to 192.168.32.1 is directed toward 10.1.1.2, because 192.168.32.1 falls within the 192.168.32.0/24 network. It also falls within the other route in the routing table, but the 192.168.32.0/24 has the longest prefix within the routing table (24 bits versus 19 bits). Longer prefixes are always preferred over shorter ones when forwarding a packet.

Dynamic Routing and Failover

Dynamic routes are not replicated to the standby unit or failover group in a failover configuration. Therefore, immediately after a failover occurs, some packets received by the security appliance may be dropped because of a lack of routing information or routed to a default static route while the routing table is repopulated by the configured dynamic routing protocols.



CHAPTER 10

Configuring DHCP, DDNS, and WCCP Services

This chapter describes how to configure the DHCP server, dynamic DNS (DDNS) update methods, and WCCP on the security appliance. DHCP provides network configuration parameters, such as IP addresses, to DHCP clients. The security appliance can provide a DHCP server or DHCP relay services to DHCP clients attached to security appliance interfaces. The DHCP server provides network configuration parameters directly to DHCP clients. DHCP relay passes DHCP requests received on one interface to an external DHCP server located behind a different interface.

DDNS update integrates DNS with DHCP. The two protocols are complementary: DHCP centralizes and automates IP address allocation; DDNS update automatically records the association between assigned addresses and hostnames at pre-defined intervals. DDNS allows frequently changing address-hostname associations to be updated frequently. Mobile hosts, for example, can then move freely on a network without user or administrator intervention. DDNS provides the necessary dynamic updating and synchronizing of the name to address and address to name mappings on the DNS server.

WCCP specifies interactions between one or more routers, Layer 3 switches, or security appliances and one or more web caches. The feature transparently redirects selected types of traffic to a group of web cache engines to optimize resource usage and lower response times.

This chapter includes the following sections:

- [Configuring a DHCP Server, page 10-1](#)
- [Configuring DHCP Relay Services, page 10-5](#)
- [Configuring Dynamic DNS, page 10-6](#)
- [Configuring Web Cache Services Using WCCP, page 10-9](#)

Configuring a DHCP Server

This section describes how to configure DHCP server provided by the security appliance. This section includes the following topics:

- [Enabling the DHCP Server, page 10-2](#)
- [Configuring DHCP Options, page 10-3](#)
- [Using Cisco IP Phones with a DHCP Server, page 10-4](#)

Enabling the DHCP Server

The security appliance can act as a DHCP server. DHCP is a protocol that supplies network settings to hosts including the host IP address, the default gateway, and a DNS server.



Note

The security appliance DHCP server does not support BOOTP requests.

In multiple context mode, you cannot enable the DHCP server or DHCP relay on an interface that is used by more than one context.

You can configure a DHCP server on each interface of the security appliance. Each interface can have its own pool of addresses to draw from. 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 services 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.

To enable the DHCP server on a given security appliance interface, perform the following steps:

- Step 1** Create a DHCP address pool. Enter the following command to define the address pool:

```
hostname(config)# dhcpd address ip_address-ip_address interface_name
```

The security appliance assigns a client one of the addresses from this pool to use for a given length of time. These addresses are the local, untranslated addresses for the directly connected network.

The address pool must be on the same subnet as the security appliance interface.

- Step 2** (Optional) To specify the IP address(es) of the DNS server(s) the client will use, enter the following command:

```
hostname(config)# dhcpd dns dns1 [dns2]
```

You can specify up to two DNS servers.

- Step 3** (Optional) To specify the IP address(es) of the WINS server(s) the client will use, enter the following command:

```
hostname(config)# dhcpd wins wins1 [wins2]
```

You can specify up to two WINS servers.

- Step 4** (Optional) To change the lease length to be granted to the client, enter the following command:

```
hostname(config)# dhcpd lease lease_length
```

This lease equals the amount of time (in seconds) the client can use its allocated IP address before the lease expires. Enter a value between 300 to 1,048,575. The default value is 3600 seconds.

- Step 5** (Optional) To configure the domain name the client uses, enter the following command:

```
hostname(config)# dhcpd domain domain_name
```

- Step 6** (Optional) To configure the DHCP ping timeout value, enter the following command:

```
hostname(config)# dhcpd ping_timeout milliseconds
```

To avoid address conflicts, the security appliance sends two ICMP ping packets to an address before assigning that address to a DHCP client. This command specifies the timeout value for those packets.

- Step 7** (Transparent Firewall Mode) Define a default gateway. To define the default gateway that is sent to DHCP clients, enter the following command.

```
hostname(config)# dhcpd option 3 ip gateway_ip
```

If you do not use the DHCP option 3 to define the default gateway, DHCP clients use the IP address of the management interface. The management interface does not route traffic.

- Step 8** To enable the DHCP daemon within the security appliance to listen for DHCP client requests on the enabled interface, enter the following command:

```
hostname(config)# dhcpd enable interface_name
```

For example, to assign the range 10.0.1.101 to 10.0.1.110 to hosts connected to the inside interface, enter the following commands:

```
hostname(config)# dhcpd address 10.0.1.101-10.0.1.110 inside
hostname(config)# dhcpd dns 209.165.201.2 209.165.202.129
hostname(config)# dhcpd wins 209.165.201.5
hostname(config)# dhcpd lease 3000
hostname(config)# dhcpd domain example.com
hostname(config)# dhcpd enable inside
```

Configuring DHCP Options

You can configure the security appliance to send information for the DHCP options listed in RFC 2132. The DHCP options fall into one of three categories:

- Options that return an IP address.
- Options that return a text string.
- Options that return a hexadecimal value.

The security appliance supports all three categories of DHCP options. To configure a DHCP option, do one of the following:

- To configure a DHCP option that returns one or two IP addresses, enter the following command:

```
hostname(config)# dhcpd option code ip addr_1 [addr_2]
```

- To configure a DHCP option that returns a text string, enter the following command:

```
hostname(config)# dhcpd option code ascii text
```

- To configure a DHCP option that returns a hexadecimal value, enter the following command:

```
hostname(config)# dhcpd option code hex value
```



Note

The security appliance does not verify that the option type and value that you provide match the expected type and value for the option code as defined in RFC 2132. For example, you can enter the **dhcpd option 46 ascii hello** command and the security appliance accepts the configuration although option 46 is defined in RFC 2132 as expecting a single-digit, hexadecimal value. For more information about the option codes and their associated types and expected values, refer to RFC 2132.

Table 10-1 shows the DHCP options that are not supported by the **dhcpd option** command.

Table 10-1 *Unsupported DHCP Options*

Option Code	Description
0	DHCPOPT_PAD
1	HCPOPT_SUBNET_MASK
12	DHCPOPT_HOST_NAME
50	DHCPOPT_REQUESTED_ADDRESS
51	DHCPOPT_LEASE_TIME
52	DHCPOPT_OPTION_OVERLOAD
53	DHCPOPT_MESSAGE_TYPE
54	DHCPOPT_SERVER_IDENTIFIER
58	DHCPOPT_RENEWAL_TIME
59	DHCPOPT_REBINDING_TIME
61	DHCPOPT_CLIENT_IDENTIFIER
67	DHCPOPT_BOOT_FILE_NAME
82	DHCPOPT_RELAY_INFORMATION
255	DHCPOPT_END

Specific options, DHCP option 3, 66, and 150, are used to configure Cisco IP Phones. See the [“Using Cisco IP Phones with a DHCP Server” section on page 10-4](#) topic for more information about configuring those options.

Using Cisco IP Phones with a DHCP Server

Enterprises with small branch offices that implement a Cisco IP Telephony Voice over IP solution typically implement Cisco CallManager at a central office to control Cisco IP Phones at small branch offices. This implementation allows centralized call processing, reduces the equipment required, and eliminates the administration of additional Cisco CallManager and other servers at branch offices.

Cisco IP Phones download their configuration from a TFTP server. When a Cisco IP Phone starts, if it does not have both the IP address and TFTP server IP address preconfigured, it sends a request with option 150 or 66 to the DHCP server to obtain this information.

- DHCP option 150 provides the IP addresses of a list of TFTP servers.
- DHCP option 66 gives the IP address or the hostname of a single TFTP server.

Cisco IP Phones might also include DHCP option 3 in their requests, which sets the default route.

Cisco IP Phones might include both option 150 and 66 in a single request. In this case, the security appliance DHCP server provides values for both options in the response if they are configured on the security appliance.

You can configure the security appliance to send information for most options listed in RFC 2132. The following example shows the syntax for any option number, as well as the syntax for commonly-used options 66, 150, and 3:

- To provide information for DHCP requests that include an option number as specified in RFC-2132, enter the following command:

```
hostname(config)# dhcpd option number value
```

- To provide the IP address or name of a TFTP server for option 66, enter the following command:

```
hostname(config)# dhcpd option 66 ascii server_name
```

- To provide the IP address or names of one or two TFTP servers for option 150, enter the following command:

```
hostname(config)# dhcpd option 150 ip server_ip1 [server_ip2]
```

The *server_ip1* is the IP address or name of the primary TFTP server while *server_ip2* is the IP address or name of the secondary TFTP server. A maximum of two TFTP servers can be identified using option 150.

- To set the default route, enter the following command:

```
hostname(config)# dhcpd option 3 ip router_ip1
```

Configuring DHCP Relay Services

A DHCP relay agent allows the security appliance to forward DHCP requests from clients to a router connected to a different interface.

The following restrictions apply to the use of the DHCP relay agent:

- The relay agent cannot be enabled if the DHCP server feature is also enabled.
- Clients must be directly connected to the security appliance and cannot send requests through another relay agent or a router.
- For multiple context mode, you cannot enable DHCP relay on an interface that is used by more than one context.



Note

DHCP Relay services are not available in transparent firewall mode. A security appliance in transparent firewall mode only allows ARP traffic through; all other traffic requires an access list. To allow DHCP requests and replies through the security appliance 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.



Note

When DHCP relay is enabled and more than one DHCP relay server is defined, the security appliance 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 security appliance receives any of the following DHCP messages: ACK, NACK, or decline.

To enable DHCP relay, perform the following steps:

- Step 1** To set the IP address of a DHCP server on a different interface from the DHCP client, enter the following command:

```
hostname(config)# dhcprelay server ip_address if_name
```

You can use this command up to 4 times to identify up to 4 servers.

- Step 2** To enable DHCP relay on the interface connected to the clients, enter the following command:

```
hostname(config)# dhcprelay enable interface
```

- Step 3** (Optional) To set the number of seconds allowed for relay address negotiation, enter the following command:

```
hostname(config)# dhcprelay timeout seconds
```

- Step 4** (Optional) To change the first default router address in the packet sent from the DHCP server to the address of the security appliance interface, enter the following command:

```
hostname(config)# dhcprelay setroute interface_name
```

This action allows the client to set its default route to point to the security appliance even if the DHCP server specifies a different router.

If there is no default router option in the packet, the security appliance adds one containing the interface address.

The following example enables the security appliance to forward DHCP requests from clients connected to the inside interface to a DHCP server on the outside interface:

```
hostname(config)# dhcprelay server 201.168.200.4
hostname(config)# dhcprelay enable inside
hostname(config)# dhcprelay setroute inside
```

Configuring Dynamic DNS

This section describes examples for configuring the security appliance to support Dynamic DNS. DDNS update integrates DNS with DHCP. The two protocols are complementary—DHCP centralizes and automates IP address allocation, while dynamic DNS update automatically records the association between assigned addresses and hostnames. When you use DHCP and dynamic DNS update, this configures a host automatically for network access whenever it attaches to the IP network. You can locate and reach the host using its permanent, unique DNS hostname. Mobile hosts, for example, can move freely without user or administrator intervention.

DDNS provides address and domain name mappings so hosts can find each other even though their DHCP-assigned IP addresses change frequently. The DDNS name and address mappings are held on the DHCP server in two resource records: the A RR contains the name to IP address mapping while the PTR RR maps addresses to names. Of the two methods for performing DDNS updates—the IETF standard defined by RFC 2136 and a generic HTTP method—the security appliance supports the IETF method in this release.

The two most common DDNS update configurations are:

- The DHCP client updates the A RR while the DHCP server updates PTR RR.
- The DHCP server updates both the A and PTR RRs.

In general, the DHCP server maintains DNS PTR RRs on behalf of clients. Clients may be configured to perform all desired DNS updates. The server may be configured to honor these updates or not. To update the PTR RR, the DHCP server must know the Fully Qualified Domain Name of the client. The client provides an FQDN to the server using a DHCP option called Client FQDN.

The following examples present these common scenarios:

- [Example 1: Client Updates Both A and PTR RRs for Static IP Addresses, page 10-7](#)

- [Example 2: Client Updates Both A and PTR RRs; DHCP Server Honors Client Update Request; FQDN Provided Through Configuration, page 10-7](#)
- [Example 3: Client Includes FQDN Option Instructing Server Not to Update Either RR; Server Overrides Client and Updates Both RRs., page 10-8](#)
- [Example 4: Client Asks Server To Perform Both Updates; Server Configured to Update PTR RR Only; Honors Client Request and Updates Both A and PTR RR, page 10-8](#)
- [Example 5: Client Updates A RR; Server Updates PTR RR, page 10-9](#)

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

The following example configures the client to request that it update both A and PTR resource records for static IP addresses. To configure this example, perform the following steps:

-
- Step 1** To define a DDNS update method called ddns-2 that requests that the client update both the A and PTR RRs, enter the following commands:
- ```
hostname(config)# ddns update method ddns-2
hostname(DDNS-update-method)# ddns both
```
- Step 2** To associate the method ddns-2 with the eth1 interface, enter the following commands:
- ```
hostname(DDNS-update-method)# interface eth1
hostname(config-if)# ddns update ddns-2
hostname(config-if)# ddns update hostname asa.example.com
```
- Step 3** To configure a static IP address for eth1, enter the following commands:
- ```
hostname(config-if)# ip address 10.0.0.40 255.255.255.0
```
- 

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

The following example configures 1) the DHCP client to request that it update both the A and PTR RRs, and 2) the DHCP server to honor the requests. To configure this example, perform the following steps:

- 
- Step 1** To configure the DHCP client to request that the DHCP server perform no updates, enter the following command:
- ```
hostname(config)# dhcp-client update dns server none
```
- Step 2** To create a DDNS update method named ddns-2 on the DHCP client that requests that the client perform both A and PTR updates, enter the following commands:
- ```
hostname(config)# ddns update method ddns-2
hostname(DDNS-update-method)# ddns both
```
- Step 3** To associate the method named ddns-2 with the security appliance interface named Ethernet0, and enable DHCP on the interface, enter the following commands:
- ```
hostname(DDNS-update-method)# interface Ethernet0
hostname(if-config)# ddns update ddns-2
hostname(if-config)# ddns update hostname asa.example.com
hostname(if-config)# ip address dhcp
```

- Step 4** To configure the DHCP server, enter the following command:

```
hostname(if-config)# dhcpd update dns
```

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

The following example configures the DHCP client to include the FQDN option instructing the DHCP server not to update either the A or PTR updates. The example also configures the server to override the client request. As a result, the client backs off without performing any updates.

To configure this scenario, perform the following steps:

- Step 1** To configure the update method named ddns-2 to request that it make both A and PTR RR updates, enter the following commands:

```
hostname(config)# ddns update method ddns-2
hostname(DDNS-update-method)# ddns both
```

- Step 2** To assign the DDNS update method named ddns-2 on interface Ethernet0 and provide the client hostname (asa), enter the following commands:

```
hostname(DDNS-update-method)# interface Ethernet0
hostname(if-config)# ddns update ddns-2
hostname(if-config)# ddns update hostname asa.example.com
```

- Step 3** To enable the DHCP client feature on the interface, enter the following commands:

```
hostname(if-config)# dhcp client update dns server none
hostname(if-config)# ip address dhcp
```

- Step 4** To configure the DHCP server to override the client update requests, enter the following command:

```
hostname(if-config)# dhcpd update dns both override
```

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

The following example configures the server to perform only PTR RR updates by default. However, the server honors the client request that it perform both A and PTR updates. The server also forms the FQDN by appending the domain name (example.com) to the hostname provided by the client (asa).

To configure this scenario, perform the following steps:

- Step 1** To configure the DHCP client on interface Ethernet0, enter the following commands:

```
hostname(config)# interface Ethernet0
hostname(config-if)# dhcp client update dns both
hostname(config-if)# ddns update hostname asa
```

- Step 2** To configure the DHCP server, enter the following commands:

```
hostname(config-if)# dhcpd update dns
```



```
hostname(config-if)# dhcpd domain example.com
```

Example 5: Client Updates A RR; Server Updates PTR RR

The following example configures the client to update the A resource record and the server to update the PTR records. Also, the client uses the domain name from the DHCP server to form the FQDN.

To configure this scenario, perform the following steps:

Step 1 To define the DDNS update method named ddns-2, enter the following commands:

```
hostname(config)# ddns update method ddns-2
hostname(DDNS-update-method)# ddns
```

Step 2 To configure the DHCP client for interface Ethernet0 and assign the update method to the interface, enter the following commands:

```
hostname(DDNS-update-method)# interface Ethernet0
hostname(config-if)# dhcp client update dns
hostname(config-if)# ddns update ddns-2
hostname(config-if)# ddns update hostname asa
```

Step 3 To configure the DHCP server, enter the following commands:

```
hostname(config-if)# dhcpd update dns
hostname(config-if)# dhcpd domain example.com
```

Configuring Web Cache Services Using WCCP

The purpose of web caching is to reduce latency and network traffic. Previously-accessed web pages are stored in a cache buffer, so if a user needs the page again, they can retrieve it from the cache instead of the web server.

WCCP specifies interactions between the security appliance and external web caches. The feature transparently redirects selected types of traffic to a group of web cache engines to optimize resource usage and lower response times. The security appliance only supports WCCP version 2.

Using a security appliance as an intermediary eliminates the need for a separate router to do the WCCP redirect because the security appliance takes care of redirecting requests to cache engines. When the security appliance knows when a packet needs redirection, it skips TCP state tracking, TCP sequence number randomization, and NAT on these traffic flows.

This section includes the following topics:

- [WCCP Feature Support, page 10-9](#)
- [WCCP Interaction With Other Features, page 10-10](#)
- [Enabling WCCP Redirection, page 10-10](#)

WCCP Feature Support

The following WCCPv2 features are supported with the security appliance:

- Redirection of multiple TCP/UDP port-destined traffic.
- Authentication for cache engines in a service group.

The following WCCPv2 features are not supported with the security appliance:

- Multiple routers in a service group is not supported. Multiple Cache Engines in a service group is still supported.
- Multicast WCCP is not supported.
- The Layer 2 redirect method is not supported; only GRE encapsulation is supported.
- WCCP source address spoofing.

WCCP Interaction With Other Features

In the security appliance implementation of WCCP, the following applies as to how the protocol interacts with other configurable features:

- An ingress access list entry always takes higher priority over WCCP. For example, if an access list does not permit a client to communicate with a server then traffic will not be redirected to a cache engine. Both ingress interface access lists and egress interface access lists will be applied.
- TCP intercept, authorization, URL filtering, inspect engines, and IPS features are not applied to a redirected flow of traffic.
- When a cache engine cannot service a request and packet is returned, or when a cache miss happens on a cache engine and it requests data from a web server, then the contents of the traffic flow will be subject to all the other configured features of the security appliance.
- In failover, WCCP redirect tables are not replicated to standby units. After a failover, packets will not be redirected until the tables are rebuilt. Sessions redirected prior to failover will likely be reset by the web server.

Enabling WCCP Redirection

There are two steps to configuring WCCP redirection on the security appliance. The first involves identifying the service to be redirected with the **wccp** command, and the second is defining on which interface the redirection occurs with the **wccp redirect** command. The **wccp** command can optionally also define which cache engines can participate in the service group, and what traffic should be redirected to the cache engine.

WCCP redirect is supported only on the ingress of an interface. The only topology that the security appliance supports is when client and cache engine are behind the same interface of the security appliance and the cache engine can directly communicate with the client without going through the security appliance.

The following configuration tasks assume you have already installed and configured the cache engines you wish to include in your network.

To configure WCCP redirection, perform the following steps:

Step 1 To enable a WCCP service group, enter the following command:

```
hostname(config)# wccp {web-cache | service_number} [redirect-list access_list]
[group-list access_list] [password password]
```

The standard service is **web-cache**, which intercepts TCP port 80 (HTTP) traffic and redirects that traffic to the cache engines, but you can identify a service number if desired between 0 and 254. For example, to transparently redirect native FTP traffic to a cache engine, use WCCP service 60. You can enter this command multiple times for each service group you want to enable.

The **redirect-list** *access_list* argument controls traffic redirected to this service group.

The **group-list** *access_list* argument determines which web cache IP addresses are allowed to participate in the service group.

The **password** *password* argument specifies MD5 authentication for messages received from the service group. Messages that are not accepted by the authentication are discarded.

Step 2 To enable WCCP redirection on an interface, enter the following command:

```
hostname(config)# wccp interface interface_name {web-cache | service_number} redirect in
```

The standard service is **web-cache**, which intercepts TCP port 80 (HTTP) traffic and redirects that traffic to the cache engines, but you can identify a service number if desired between 0 and 254. For example, to transparently redirect native FTP traffic to a cache engine, use WCCP service 60. You can enter this command multiple times for each service group you want to participate in.

For example, to enable the standard **web-cache** service and redirect HTTP traffic that enters the inside interface to a web cache, enter the following commands:

```
hostname(config)# wccp web-cache
hostname(config)# wccp interface inside web-cache redirect in
```




CHAPTER 11

Configuring Multicast Routing

This chapter describes how to configure multicast routing. This section includes the following topics:

- [Multicast Routing Overview, page 11-13](#)
- [Enabling Multicast Routing, page 11-14](#)
- [Configuring IGMP Features, page 11-14](#)
- [Configuring Stub Multicast Routing, page 11-17](#)
- [Configuring a Static Multicast Route, page 11-17](#)
- [Configuring PIM Features, page 11-18](#)
- [For More Information about Multicast Routing, page 11-22](#)

Multicast Routing Overview

The security appliance supports both stub multicast routing and PIM multicast routing. However, you cannot configure both concurrently on a single security appliance.

Stub multicast routing provides dynamic host registration and facilitates multicast routing. When configured for stub multicast routing, the security appliance acts as an IGMP proxy agent. Instead of fully participating in multicast routing, the security appliance forwards IGMP messages to an upstream multicast router, which sets up delivery of the multicast data. When configured for stub multicast routing, the security appliance cannot be configured for PIM.

The security appliance supports both PIM-SM and bi-directional 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.

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 security appliance is the PIM RP, use the untranslated outside address of the security appliance as the RP address.

Enabling Multicast Routing

Enabling multicast routing lets the security appliance forward multicast packets. Enabling multicast routing automatically enables PIM and IGMP on all interfaces. To enable multicast routing, enter the following command:

```
hostname(config)# multicast-routing
```

The number of entries in the multicast routing tables are limited by the amount of RAM on the system. [Table 11-1](#) lists the maximum number of entries for specific multicast tables based on the amount of RAM on the security appliance. Once these limits are reached, any new entries are discarded.

Table 11-1 *Entry Limits for Multicast Tables*

Table	16 MB	128 MB	128+ MB
MFIB	1000	3000	5000
IGMP Groups	1000	3000	5000
PIM Routes	3000	7000	12000

Configuring IGMP Features

IP hosts use IGMP to report their group memberships to directly connected multicast routers. IGMP uses group addresses (Class D IP address) as group identifiers. Host group address can be in the range 224.0.0.0 to 239.255.255.255. The address 224.0.0.0 is never assigned to any group. The address 224.0.0.1 is assigned to all systems on a subnet. The address 224.0.0.2 is assigned to all routers on a subnet.

When you enable multicast routing on the security appliance, 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. This section includes the following topics:

- [Disabling IGMP on an Interface, page 11-15](#)
- [Configuring Group Membership, page 11-15](#)
- [Configuring a Statically Joined Group, page 11-15](#)
- [Controlling Access to Multicast Groups, page 11-15](#)
- [Limiting the Number of IGMP States on an Interface, page 11-16](#)
- [Modifying the Query Interval and Query Timeout, page 11-16](#)
- [Changing the Query Response Time, page 11-17](#)
- [Changing the IGMP Version, page 11-17](#)

Disabling IGMP on an Interface

You can disable IGMP on specific interfaces. This is useful if you know that you do not have any multicast hosts on a specific interface and you want to prevent the security appliance from sending host query messages on that interface.

To disable IGMP on an interface, enter the following command:

```
hostname(config-if)# no igmp
```

To reenale IGMP on an interface, enter the following command:

```
hostname(config-if)# igmp
```

**Note**

Only the **no igmp** command appears in the interface configuration.

Configuring Group Membership

You can configure the security appliance to be a member of a multicast group. Configuring the security appliance 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.

To have the security appliance join a multicast group, enter the following command:

```
hostname(config-if)# igmp join-group group-address
```

Configuring a Statically Joined Group

Sometimes a group member cannot report its membership in the group, or there may be no members of a group on the network segment, but 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 in one of two ways:

- Using the **igmp join-group** command (see [Configuring Group Membership, page 11-15](#)). This causes the security appliance to accept and to forward the multicast packets.
- Using the **igmp static-group** command. The security appliance does not accept the multicast packets but rather forwards them to the specified interface.

To configure a statically joined multicast group on an interface, enter the following command:

```
hostname(config-if)# igmp static-group group-address
```

Controlling Access to Multicast Groups

To control the multicast groups that hosts on the security appliance interface can join, perform the following steps:

Step 1 Create an access list for the multicast traffic. You can create more than one entry for a single access list. You can use extended or standard access lists.

- To create a standard access list, enter the following command:

```
hostname(config)# access-list name standard [permit | deny] ip_addr mask
```

The *ip_addr* argument is the IP address of the multicast group being permitted or denied.

- To create an extended access list, enter the following command:

```
hostname(config)# access-list name extended [permit | deny] protocol src_ip_addr  
src_mask dst_ip_addr dst_mask
```

The *dst_ip_addr* argument is the IP address of the multicast group being permitted or denied.

Step 2 Apply the access list to an interface by entering the following command:

```
hostname(config-if)# igmp access-group acl
```

The *acl* argument is the name of a standard or extended IP access list.

Limiting the Number of IGMP States on an Interface

You can limit the number of IGMP states resulting from IGMP membership reports on a per-interface basis. Membership reports exceeding the configured limits are not entered in the IGMP cache and traffic for the excess membership reports is not forwarded.

To limit the number of IGMP states on an interface, enter the following command:

```
hostname(config-if)# igmp limit number
```

Valid values range from 0 to 500, with 500 being the default value. Setting this value to 0 prevents learned groups from being added, but manually defined memberships (using the **igmp join-group** and **igmp static-group** commands) are still permitted. The **no** form of this command restores the default value.

Modifying the Query Interval and Query Timeout

The security appliance 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 security appliance. If the security appliance 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. To change this interval, enter the following command:

```
hostname(config-if)# igmp query-interval seconds
```

If the security appliance does not hear a query message on an interface for the specified timeout value (by default, 255 seconds), then the security appliance becomes the designated router and starts sending the query messages. To change this timeout value, enter the following command:

```
hostname(config-if)# igmp query-timeout seconds
```


**Note**

The **igmp query-timeout** and **igmp query-interval** commands require IGMP Version 2.

Changing the Query Response Time

By default, the maximum query response time advertised in IGMP queries is 10 seconds. If the security appliance does not receive a response to a host query within this amount of time, it deletes the group.

To change the maximum query response time, enter the following command:

```
hostname(config-if)# igmp query-max-response-time seconds
```

Changing the IGMP Version

By default, the security appliance runs IGMP Version 2, which enables several additional features such as the **igmp query-timeout** and **igmp query-interval** commands.

All multicast routers on a subnet must support the same version of IGMP. The security appliance 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 security appliance running IGMP Version 2 works correctly when IGMP Version 1 hosts are present.

To control which version of IGMP is running on an interface, enter the following command:

```
hostname(config-if)# igmp version {1 | 2}
```

Configuring Stub Multicast Routing

A security appliance 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. To configure the security appliance as an IGMP proxy agent, forward the host join and leave messages from the stub area interface to an upstream interface.

To forward the host join and leave messages, enter the following command from the interface attached to the stub area:

```
hostname(config-if)# igmp forward interface if_name
```

**Note**

Stub Multicast Routing and PIM are not supported concurrently.

Configuring a Static Multicast Route

When using PIM, the security appliance expects to receive packets on the same interface where it sends unicast packets back to the source. In some cases, such as bypassing a route that does not support multicast routing, you may want unicast packets to take one path and multicast packets to take another.

Static multicast routes are not advertised or redistributed.

To configure a static multicast route for PIM, enter the following command:

```
hostname(config)# mroute src_ip src_mask {input_if_name | rpf_addr} [distance]
```

To configure a static multicast route for a stub area, enter the following command:

```
hostname(config)# mroute src_ip src_mask input_if_name [dense output_if_name] [distance]
```

**Note**

The **dense output_if_name** keyword and argument pair is only supported for stub multicast routing.

Configuring PIM Features

Routers use PIM to maintain forwarding tables for forwarding multicast diagrams. When you enable multicast routing on the security appliance, 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. This section includes the following topics:

- [Disabling PIM on an Interface, page 11-18](#)
- [Configuring a Static Rendezvous Point Address, page 11-19](#)
- [Configuring the Designated Router Priority, page 11-19](#)
- [Filtering PIM Register Messages, page 11-19](#)
- [Configuring PIM Message Intervals, page 11-20](#)
- [Configuring a Multicast Boundary, page 11-20](#)
- [Filtering PIM Neighbors, page 11-20](#)
- [Supporting Mixed Bidirectional/Sparse-Mode PIM Networks, page 11-21](#)

Disabling PIM on an Interface

You can disable PIM on specific interfaces. To disable PIM on an interface, enter the following command:

```
hostname(config-if)# no pim
```

To reenable PIM on an interface, enter the following command:

```
hostname(config-if)# pim
```

**Note**

Only the **no pim** command appears in the interface configuration.

Configuring a Static Rendezvous Point Address

All routers within a common PIM sparse mode or bidir domain require knowledge of the PIM RP address. The address is statically configured using the **pim rp-address** command.

**Note**

The security appliance does not support Auto-RP or PIM BSR; you must use the **pim rp-address** command to specify the RP address.

You can configure the security appliance to serve as RP to more than one group. The group range specified in the access list determines the PIM RP group mapping. If an access list is not specified, then the RP for the group is applied to the entire multicast group range (224.0.0.0/4).

To configure the address of the PIM RP, enter the following command:

```
hostname(config)# pim rp-address ip_address [acl] [bidir]
```

The *ip_address* argument is the unicast IP address of the router to be a PIM RP. The *acl* argument is the name or number of a standard access list that defines which multicast groups the RP should be used with. Do not use a host ACL with this command. Excluding the **bidir** keyword causes the groups to operate in PIM sparse mode.

**Note**

The security appliance always advertises the bidir capability in the PIM hello messages regardless of the actual bidir configuration.

Configuring the Designated Router Priority

The DR is responsible for sending PIM register, join, and prune messages to the RP. When there is more than one multicast router on a network segment, there is an election process to select the DR based on DR priority. If multiple devices have the same DR priority, then the device with the highest IP address becomes the DR.

By default, the security appliance has a DR priority of 1. You can change this value by entering the following command:

```
hostname(config-if)# pim dr-priority num
```

The *num* argument can be any number from 1 to 4294967294.

Filtering PIM Register Messages

You can configure the security appliance to filter PIM register messages. To filter PIM register messages, enter the following command:

```
hostname(config)# pim accept-register {list acl | route-map map-name}
```

Configuring PIM Message Intervals

Router query messages are used to elect the PIM DR. The PIM DR is responsible for sending router query messages. By default, router query messages are sent every 30 seconds. You can change this value by entering the following command:

```
hostname(config-if)# pim hello-interval seconds
```

Valid values for the *seconds* argument range from 1 to 3600 seconds.

Every 60 seconds, the security appliance sends PIM join/prune messages. To change this value, enter the following command:

```
hostname(config-if)# pim join-prune-interval seconds
```

Valid values for the *seconds* argument range from 10 to 600 seconds.

Configuring a Multicast Boundary

Address scoping defines domain boundaries so that domains with RPs that have the same IP address do not leak into each other. Scoping is performed on the subnet boundaries within large domains and on the boundaries between the domain and the Internet.

You can set up an administratively scoped boundary on an interface for multicast group addresses using the **multicast boundary** command. IANA has designated the multicast address range 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. They would be considered local, not globally unique.

To configure a multicast boundary, enter the following command:

```
hostname(config-if)# multicast boundary acl [filter-autorp]
```

A standard ACL defines the range of addresses affected. 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 the **filter-autorp** keyword to 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 access control list (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.

Filtering PIM Neighbors

You can define the routers that can become PIM neighbors with the **pim neighbor-filter** command. By filtering the routers that can become PIM neighbors, you can:

- Prevent unauthorized routers from becoming PIM neighbors.
- Prevent attached stub routers from participating in PIM.

To define the neighbors that can become a PIM neighbor, perform the following steps:

-
- Step 1** Use the **access-list** command to define a standard access list defines the routers you want to participate in PIM.

For example the following access list, when used with the **pim neighbor-filter** command, prevents the 10.1.1.1 router from becoming a PIM neighbor:

```
hostname(config)# access-list pim_nbr deny 10.1.1.1 255.255.255.255
```

- Step 2** Use the **pim neighbor-filter** command on an interface to filter the neighbor routers.

For example, the following commands prevent the 10.1.1.1 router from becoming a PIM neighbor on interface GigabitEthernet0/3:

```
hostname(config)# interface GigabitEthernet0/3
hostname(config-if)# pim neighbor-filter pim_nbr
```

Supporting Mixed Bidirectional/Sparse-Mode PIM Networks

Bidirectional PIM allows multicast routers to keep reduced state information. All of the multicast routers in a segment must be bidirectionally enabled in order for bidir to elect a DF.

The **pim bidir-neighbor-filter** command enables the transition from a sparse-mode-only network to a bidir network by letting you specify the routers that should participate in DF election while still allowing all routers to participate in the sparse-mode domain. The bidir-enabled routers can elect a DF from among themselves, even when there are non-bidir routers on the segment. Multicast boundaries on the non-bidir routers prevent PIM messages and data from the bidir groups from leaking in or out of the bidir subset cloud.

When the **pim bidir-neighbor-filter** command is enabled, the routers that are permitted by the ACL are considered to be bidir-capable. Therefore:

- If a permitted neighbor does not support bidir, the DF election does not occur.
- If a denied neighbor supports bidir, then DF election does not occur.
- If a denied neighbor does not support bidir, the DF election occurs.

To control which neighbors can participate in the DF election, perform the following steps:

-
- Step 1** Use the **access-list** command to define a standard access list that permits the routers you want to participate in the DF election and denies all others.

For example, the following access list permits the routers at 10.1.1.1 and 10.2.2.2 to participate in the DF election and denies all others:

```
hostname(config)# access-list pim_bidir permit 10.1.1.1 255.255.255.255
hostname(config)# access-list pim_bidir permit 10.1.1.2 255.255.255.255
hostname(config)# access-list pim_bidir deny any
```

- Step 2** Enable the **pim bidir-neighbor-filter** command on an interface.

The following example applies the access list created previous step to the interface GigabitEthernet0/3.

```
hostname(config)# interface GigabitEthernet0/3
hostname(config-if)# pim bidir-neighbor-filter pim_bidir
```

For More Information about Multicast Routing

The following RFCs from the IETF provide technical details about the IGMP and multicast routing standards used for implementing the SMR feature:

- RFC 2236 IGMPv2
- RFC 2362 PIM-SM
- RFC 2588 IP Multicast and Firewalls
- RFC 2113 IP Router Alert Option
- IETF draft-ietf-idmr-igmp-proxy-01.txt



CHAPTER 12

Configuring IPv6

This chapter describes how to enable and configure IPv6 on the security appliance. IPv6 is available in Routed firewall mode only.

This chapter includes the following sections:

- [IPv6-enabled Commands, page 12-1](#)
- [Configuring IPv6, page 12-2](#)
- [Verifying the IPv6 Configuration, page 12-11](#)

For an sample IPv6 configuration, see [Appendix B, “Sample Configurations.”](#)

IPv6-enabled Commands

The following security appliance commands can accept and display IPv6 addresses:

- capture
- configure
- copy
- http
- name
- object-group
- ping
- show conn
- show local-host
- show tcpstat
- ssh
- telnet
- tftp-server
- who
- write

**Note**

Failover does not support IPv6. The **ipv6 address** command does not support setting standby addresses for failover configurations. The **failover interface ip** command does not support using IPv6 addresses on the failover and Stateful Failover interfaces.

When entering IPv6 addresses in commands that support them, simply enter the IPv6 address using standard IPv6 notation, for example `ping fe80::2e0:b6ff:fe01:3b7a`. The security appliance correctly recognizes and processes the IPv6 address. However, you must enclose the IPv6 address in square brackets ([]) in the following situations:

- You need to specify a port number with the address, for example `[fe80::2e0:b6ff:fe01:3b7a]:8080`.
- The command uses a colon as a separator, such as the **write net** and **config net** commands, for example `configure net [fe80::2e0:b6ff:fe01:3b7a]:/tftp/config/pixconfig`.

The following commands were modified to work for IPv6:

- debug
- fragment
- ip verify
- mtu
- icmp (entered as **ipv6 icmp**)

The following inspection engines support IPv6:

- FTP
- HTTP
- ICMP
- SMTP
- TCP
- UDP

Configuring IPv6

This section contains the following topics:

- [Configuring IPv6 on an Interface, page 12-3](#)
- [Configuring a Dual IP Stack on an Interface, page 12-4](#)
- [Enforcing the Use of Modified EUI-64 Interface IDs in IPv6 Addresses, page 12-4](#)
- [Configuring IPv6 Duplicate Address Detection, page 12-4](#)
- [Configuring IPv6 Default and Static Routes, page 12-5](#)
- [Configuring IPv6 Access Lists, page 12-6](#)
- [Configuring IPv6 Neighbor Discovery, page 12-7](#)
- [Configuring a Static IPv6 Neighbor, page 12-11](#)

Configuring IPv6 on an Interface

At a minimum, each interface needs to be configured with an IPv6 link-local address. Additionally, you can add a site-local and global address to the interface.



Note

The security appliance does not support IPv6 anycast addresses.

You can configure both IPv6 and IPv4 addresses on an interface.

To configure IPv6 on an interface, perform the following steps:

- Step 1** Enter interface configuration mode for the interface on which you are configuring the IPv6 addresses:

```
hostname(config)# interface if
```

- Step 2** Configure an IPv6 address on the interface. You can assign several IPv6 addresses to an interface, such as an IPv6 link-local, site-local, and global address. However, at a minimum, you must configure a link-local address.

There are several methods for configuring IPv6 addresses. Pick the method that suits your needs from the following:

- The simplest method is to enable stateless autoconfiguration on the interface. Enabling stateless autoconfiguration on the interface configures IPv6 addresses based on prefixes received in Router Advertisement messages. A link-local address, based on the Modified EUI-64 interface ID, is automatically generated for the interface when stateless autoconfiguration is enabled. To enable stateless autoconfiguration, enter the following command:
- If you only need to configure a link-local address on the interface and are not going to assign any other IPv6 addresses to the interface, you have the option of manually defining the link-local address or generating one based on the interface MAC address (Modified EUI-64 format):

```
hostname(config-if)# ipv6 address autoconfig
```

- Enter the following command to manually specify the link-local address:

```
hostname(config-if)# ipv6 address ipv6-address link-local
```

- Enter the following command to enable IPv6 on the interface and automatically generate the link-local address using the Modified EUI-64 interface ID based on the interface MAC address:

```
hostname(config-if)# ipv6 enable
```



Note

You do not need to use the **ipv6 enable** command if you enter any other **ipv6 address** commands on an interface; IPv6 support is automatically enabled as soon as you assign an IPv6 address to the interface.

- Assign a site-local or global address to the interface. When you assign a site-local or global address, a link-local address is automatically created. Enter the following command to add a global or site-local address to the interface. Use the optional **eui-64** keyword to use the Modified EUI-64 interface ID in the low order 64 bits of the address.

```
hostname(config-if)# ipv6 address ipv6-address [eui-64]
```

- Step 3** (Optional) Suppress Router Advertisement messages on an interface. By default, Router Advertisement messages are automatically sent in response to router solicitation messages. You may want to disable these messages on any interface for which you do not want the security appliance to supply the IPv6 prefix (for example, the outside interface).

Enter the following command to suppress Router Advertisement messages on an interface:

```
hostname(config-if)# ipv6 nd suppress-ra
```

Configuring a Dual IP Stack on an Interface

The security appliance supports the configuration of both IPv6 and IPv4 on an interface. You do not need to enter any special commands to do so; simply enter the IPv4 configuration commands and IPv6 configuration commands as you normally would. Make sure you configure a default route for both IPv4 and IPv6.

Enforcing the Use of Modified EUI-64 Interface IDs in IPv6 Addresses

RFC 3513: Internet Protocol Version 6 (IPv6) Addressing Architecture requires that the interface identifier portion of all unicast IPv6 addresses, except those that start with binary value 000, be 64 bits long and be constructed in Modified EUI-64 format. The security appliance can enforce this requirement for hosts attached to the local link.

To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, enter the following command:

```
hostname(config)# ipv6 enforce-eui64 if_name
```

The *if_name* argument is the name of the interface, as specified by the **namif** command, on which you are enabling the address format enforcement.

When this command is enabled on an interface, the source addresses of IPv6 packets received on that interface are verified against the source MAC addresses to ensure that the interface identifiers use the Modified EUI-64 format. If the IPv6 packets do not use the Modified EUI-64 format for the interface identifier, the packets are dropped and the following system log message is generated:

```
%PIX|ASA-3-325003: EUI-64 source address check failed.
```

The address format verification is only performed when a flow is created. Packets from an existing flow are not checked. Additionally, the address verification can only be performed for hosts on the local link. Packets received from hosts behind a router will fail the address format verification, and be dropped, because their source MAC address will be the router MAC address and not the host MAC address.

Configuring IPv6 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:

```
%PIX|ASA-4-325002: Duplicate address ipv6_address/MAC_address on interface
```

If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface. If the duplicate address is a global address, the address is not used. However, all configuration commands associated with the duplicate address remain as configured while the state of the address is set to DUPLICATE.

If the link-local address for an interface changes, duplicate address detection is performed on the new link-local address and all of the other IPv6 address associated with the interface are regenerated (duplicate address detection is performed only on the new link-local address).

The security appliance uses neighbor solicitation messages to perform duplicate address detection. By default, the number of times an interface performs duplicate address detection is 1.

To change the number of duplicate address detection attempts, enter the following command:

```
hostname(config-if)# ipv6 nd dad attempts value
```

The *value* argument can be any value from 0 to 600. Setting the *value* argument to 0 disables duplicate address detection on the interface.

When you configure an interface to send out more than one duplicate address detection attempt, you can also use the **ipv6 nd ns-interval** command to configure the interval at which the neighbor solicitation messages are sent out. By default, they are sent out once every 1000 milliseconds.

To change the neighbor solicitation message interval, enter the following command:

```
hostname(config-if)# ipv6 nd ns-interval value
```

The *value* argument can be from 1000 to 3600000 milliseconds.

**Note**

Changing this value changes it for all neighbor solicitation messages sent out on the interface, not just those used for duplicate address detection.

Configuring IPv6 Default and Static Routes

The security appliance automatically routes IPv6 traffic between directly connected hosts if the interfaces to which the hosts are attached are enabled for IPv6 and the IPv6 ACLs allow the traffic.

The security appliance does not support dynamic routing protocols. Therefore, to route IPv6 traffic to a non-connected host or network, you need to define a static route to the host or network or, at a minimum, a default route. Without a static or default route defined, traffic to non-connected hosts or networks generate the following error message:

```
%PIX|ASA-6-110001: No route to dest_address from source_address
```

You can add a default route and static routes using the **ipv6 route** command.

To configure an IPv6 default route and static routes, perform the following steps:

Step 1 To add the default route, use the following command:

```
hostname(config)# ipv6 route if_name ::/0 next_hop_ipv6_addr
```

The address ::/0 is the IPv6 equivalent of “any.”

Step 2 (Optional) Define IPv6 static routes. Use the following command to add an IPv6 static route to the IPv6 routing table:

```
hostname(config)# ipv6 route if_name destination next_hop_ipv6_addr [admin_distance]
```



Note

The **ipv6 route** command works like the **route** command used to define IPv4 static routes.

Configuring IPv6 Access Lists

Configuring an IPv6 access list is similar configuring an IPv4 access, but with IPv6 addresses.

To configure an IPv6 access list, perform the following steps:

Step 1 Create an access entry. To create an access list, use the **ipv6 access-list** command to create entries for the access list. There are two main forms of this command to choose from, one for creating access list entries specifically for ICMP traffic, and one to create access list entries for all other types of IP traffic.

- To create an IPv6 access list entry specifically for ICMP traffic, enter the following command:

```
hostname(config)# ipv6 access-list id [line num] {permit | deny} icmp source  
destination [icmp_type]
```

- To create an IPv6 access list entry, enter the following command:

```
hostname(config)# ipv6 access-list id [line num] {permit | deny} protocol source  
[src_port] destination [dst_port]
```

The following describes the arguments for the **ipv6 access-list** command:

- id*—The name of the access list. Use the same *id* in each command when you are entering multiple entries for an access list.
- line num*—When adding an entry to an access list, you can specify the line number in the list where the entry should appear.
- permit | deny**—Determines whether the specified traffic is blocked or allowed to pass.
- icmp**—Indicates that the access list entry applies to ICMP traffic.
- protocol*—Specifies the traffic being controlled by the access list entry. This can be the name (**ip**, **tcp**, or **udp**) or number (1-254) of an IP protocol. Alternatively, you can specify a protocol object group using **object-group** *grp_id*.
- source and destination*—Specifies the source or destination of the traffic. The source or destination can be an IPv6 prefix, in the format *prefix/length*, to indicate a range of addresses, the keyword **any**, to specify any address, or a specific host designated by **host** *host_ipv6_addr*.

- *src_port and dst_port*—The source and destination port (or service) argument. Enter an operator (**lt** for less than, **gt** for greater than, **eq** for equal to, **neq** for not equal to, or **range** for an inclusive range) followed by a space and a port number (or two port numbers separated by a space for the **range** keyword).
- *icmp_type*—Specifies the ICMP message type being filtered by the access rule. The value can be a valid ICMP type number (from 0 to 155) or one of the ICMP type literals as shown in [Appendix D, “Addresses, Protocols, and Ports”](#). Alternatively, you can specify an ICMP object group using **object-group id**.

Step 2 To apply the access list to an interface, enter the following command:

```
hostname(config)# access-group access_list_name {in | out} interface if_name
```

Configuring IPv6 Neighbor Discovery

The IPv6 neighbor discovery process uses ICMPv6 messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the reachability of a neighbor, and keep track of neighboring routers.

This section contains the following topics:


- [Configuring Neighbor Solicitation Messages, page 12-7](#)
- [Configuring Router Advertisement Messages, page 12-9](#)
- [Multicast Listener Discovery Support, page 12-11](#)

Configuring Neighbor Solicitation Messages

Neighbor solicitation messages (ICMPv6 Type 135) are sent on the local link by nodes attempting to discover the link-layer addresses of other nodes on the local link. The neighbor solicitation message is sent to the solicited-node multicast address. The source address in the neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The neighbor solicitation message also includes the link-layer address of the source node.

After receiving a neighbor solicitation message, the destination node replies by sending a neighbor advertisement message (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 12-1](#) shows the neighbor solicitation and response process.

Figure 12-1 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.

You can configure the neighbor solicitation message interval and neighbor reachable time on a per-interface basis. See the following topics for more information:

- [Configuring the Neighbor Solicitation Message Interval, page 12-8](#)
- [Configuring the Neighbor Reachable Time, page 12-8](#)

Configuring the Neighbor Solicitation Message Interval

To configure the interval between IPv6 neighbor solicitation retransmissions on an interface, enter the following command:

```
hostname(config-if)# ipv6 nd ns-interval value
```

Valid values for the *value* argument range from 1000 to 3600000 milliseconds. The default value is 1000 milliseconds.

This setting is also sent in router advertisement messages.

Configuring the Neighbor Reachable Time

The neighbor reachable time enables detecting unavailable neighbors. Shorter configured times enable detecting unavailable neighbors more quickly; however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.

To configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred, enter the following command:

```
hostname(config-if)# ipv6 nd reachable-time value
```

Valid values for the *value* argument range from 0 to 3600000 milliseconds. The default is 0.


This information is also sent in router advertisement messages.

When 0 is used for the *value*, the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value. To see the time used by the security appliance when this value is set to 0, use the **show ipv6 interface** command to display information about the IPv6 interface, including the ND reachable time being used.

Configuring Router Advertisement Messages

Router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface of security appliance. The router advertisement messages are sent to the all-nodes multicast address.

Figure 12-2 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 security appliance 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. See the following topics for information about changing these settings:

- [Configuring the Router Advertisement Transmission Interval, page 12-10](#)
- [Configuring the Router Lifetime Value, page 12-10](#)
- [Configuring the IPv6 Prefix, page 12-10](#)
- [Suppressing Router Advertisement Messages, page 12-11](#)

Configuring the Router Advertisement Transmission Interval

By default, router advertisements are sent out every 200 seconds. To change the interval between router advertisement transmissions on an interface, enter the following command:

```
ipv6 nd ra-interval [msec] value
```

Valid values range from 3 to 1800 seconds (or 500 to 1800000 milliseconds if the **msec** keyword is used).

The interval between transmissions should be less than or equal to the IPv6 router advertisement lifetime if security appliance 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 desired value.

Configuring the Router Lifetime Value

The router lifetime value specifies how long nodes on the local link should consider security appliance as the default router on the link.

To configure the router lifetime value in IPv6 router advertisements on an interface, enter the following command:

```
hostname(config-if)# ipv6 nd ra-lifetime seconds
```

Valid values range from 0 to 9000 seconds. The default is 1800 seconds. Entering 0 indicates that security appliance should not be considered a default router on the selected interface.

Configuring the IPv6 Prefix

Stateless autoconfiguration uses IPv6 prefixes provided in router advertisement messages to create the global unicast address from the link-local address.

To configure which IPv6 prefixes are included in IPv6 router advertisements, enter the following command:

```
hostname(config-if)# ipv6 nd prefix ipv6-prefix/prefix-length
```



Note

For stateless autoconfiguration to work properly, the advertised prefix length in router advertisement messages must always be 64 bits.

Suppressing Router Advertisement Messages

By default, Router Advertisement messages are automatically sent in response to router solicitation messages. You may want to disable these messages on any interface for which you do not want security appliance to supply the IPv6 prefix (for example, the outside interface).

To suppress IPv6 router advertisement transmissions on an interface, enter the following command:

```
hostname(config-if)# ipv6 nd suppress-ra
```

Entering this command causes the security appliance to appear as a regular IPv6 neighbor on the link and not as an IPv6 router.

Multicast Listener Discovery Support

Multicast Listener Discovery Protocol (MLD) Version 2 is supported to discover the presence of multicast address listeners on their directly attached links, and to discover specifically which multicast addresses are of interest to those neighboring nodes. ASA becomes a multicast address listener, or a host, but not a multicast router, and responds to Multicast Listener Queries and sends Multicast Listener Reports only.

The following commands were added or enhanced to support MLD:

- `clear ipv6 mld traffic` Command
- `show ipv6 mld` Command

Configuring a Static IPv6 Neighbor

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.

To configure a static entry in the IPv6 neighbor discovery cache, enter the following command:

```
hostname(config-if)# ipv6 neighbor ipv6_address if_name mac_address
```

The *ipv6_address* argument is the link-local IPv6 address of the neighbor, the *if_name* argument is the interface through which the neighbor is available, and the *mac_address* argument is the MAC address of the neighbor interface.



Note

The **clear ipv6 neighbors** command does not remove static entries from the IPv6 neighbor discovery cache; it only clears the dynamic entries.

Verifying the IPv6 Configuration

This section describes how to verify your IPv6 configuration. You can use various clear, and show commands to verify your IPv6 settings.

This section includes the following topics:

- [The show ipv6 interface Command, page 12-12](#)

- [The show ipv6 route Command, page 12-12](#)
- [The show ipv6 mld traffic Command, page 12-13](#)

The show ipv6 interface Command

To display the IPv6 interface settings, enter the following command:

```
hostname# show ipv6 interface [if_name]
```

Including the interface name, such as “outside”, displays the settings for the specified interface. Excluding the name from the command displays the setting for all interfaces that have IPv6 enabled on them. The output for the command shows the following:

- The name and status of the interface.
- The link-local and global unicast addresses.
- The multicast groups the interface belongs to.
- ICMP redirect and error message settings.
- Neighbor discovery settings.

The following is sample output from the **show ipv6 interface** command:

```
hostname# show ipv6 interface

ipv6interface is down, line protocol is down
  IPv6 is enabled, link-local address is fe80::20d:88ff:feee:6a82 [TENTATIVE]
  No global unicast address is configured
  Joined group address(es):
    ff02::1
    ff02::1:ffee:6a82
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
```



Note

The **show interface** command only displays the IPv4 settings for an interface. To see the IPv6 configuration on an interface, you need to use the **show ipv6 interface** command. The **show ipv6 interface** command does not display any IPv4 settings for the interface (if both types of addresses are configured on the interface).

The show ipv6 route Command

To display the routes in the IPv6 routing table, enter the following command:

```
hostname# show ipv6 route
```

The output from the **show ipv6 route** command is similar to the IPv4 **show route** command. It displays the following information:

- The protocol that derived the route.
- The IPv6 prefix of the remote network.
- The administrative distance and metric for the route.
- The address of the next-hop router.

- The interface through which the next hop router to the specified network is reached.

The following is sample output from the **show ipv6 route** command:

```
hostname# show ipv6 route

IPv6 Routing Table - 7 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
L   fe80::/10 [0/0]
    via ::, inside
L   fec0::a:0:0:a0a:a70/128 [0/0]
    via ::, inside
C   fec0:0:0:a::/64 [0/0]
    via ::, inside
L   ff00::/8 [0/0]
    via ::, inside
```

The show ipv6 mld traffic Command

To display the MLD traffic counters in the IPv6 routing table, enter the following command:

```
hostname# show ipv6 mld traffic
```

The output from the **show ipv6 mld traffic** command displays whether the expected number of MLD protocol messages have been received and sent.

The following is sample output from the **show ipv6 mld traffic** command:

```
hostname# show ipv6 mld traffic
show ipv6 mld traffic
MLD Traffic Counters
Elapsed time since counters cleared: 00:01:19
      Received      Sent
Valid MLD Packets  1          3
Queries            1          0
Reports            0          3
Leaves             0          0
Mtrace packets     0          0
Errors:
Malformed Packets  0
Martian source     0
Non link-local source 0
Hop limit is not equal to 1 0
```




CHAPTER 13

Configuring AAA Servers and the Local Database

This chapter describes support for AAA (pronounced “triple A”) and how to configure AAA servers and the local database.

This chapter contains the following sections:

- [AAA Overview, page 13-1](#)
- [AAA Server and Local Database Support, page 13-2](#)
- [Configuring the Local Database, page 13-10](#)
- [Identifying AAA Server Groups and Servers, page 13-12](#)
- [Using Certificates and User Login Credentials, page 13-15](#)
- [Supporting a Zone Labs Integrity Server, page 13-16](#)

AAA Overview

AAA enables the security appliance to determine who the user is (authentication), what the user can do (authorization), and what the user did (accounting).

AAA provides an extra level of protection and control for user access than using access lists alone. For example, you can create an access list allowing all outside users to access Telnet on a server on the DMZ network. If you want only some users to access the server and you might not always know IP addresses of these users, you can enable AAA to allow only authenticated and/or authorized users to make it through the security appliance. (The Telnet server enforces authentication, too; the security appliance prevents unauthorized users from attempting to access the server.)

You can use authentication alone or with authorization and accounting. Authorization always requires a user to be authenticated first. You can use accounting alone, or with authentication and authorization.

This section includes the following topics:

- [About Authentication, page 13-1](#)
- [About Authorization, page 13-2](#)
- [About Accounting, page 13-2](#)

About Authentication

Authentication controls access by requiring valid user credentials, which are typically a username and password. You can configure the security appliance to authenticate the following items:

- All administrative connections to the security appliance including the following sessions:
 - Telnet
 - SSH
 - Serial console
 - ASDM (using HTTPS)
 - VPN management access
- The **enable** command
- Network access
- VPN access

About Authorization

Authorization controls access *per user* after users authenticate. You can configure the security appliance to authorize the following items:

- Management commands
- Network access
- VPN access

Authorization controls the services and commands available to each authenticated user. Were you not to enable authorization, authentication alone would provide the same access to services for all authenticated users.

If you need the control that authorization provides, you can configure a broad authentication rule, and then have a detailed authorization configuration. For example, you authenticate inside users who attempt to access any server on the outside network and then limit the outside servers that a particular user can access using authorization.

The security appliance caches the first 16 authorization requests per user, so if the user accesses the same services during the current authentication session, the security appliance does not resend the request to the authorization server.

About Accounting

Accounting tracks traffic that passes through the security appliance, enabling you to have a record of user activity. If you enable authentication for that traffic, you can account for traffic per user. If you do not authenticate the traffic, you can account for traffic per IP address. Accounting information includes when sessions start and stop, username, the number of bytes that pass through the security appliance for the session, the service used, and the duration of each session.

AAA Server and Local Database Support

The security appliance supports a variety of AAA server types and a local database that is stored on the security appliance. This section describes support for each AAA server type and the local database.

This section contains the following topics:

- [Summary of Support, page 13-3](#)

- [RADIUS Server Support, page 13-3](#)
- [TACACS+ Server Support, page 13-4](#)
- [SDI Server Support, page 13-4](#)
- [NT Server Support, page 13-5](#)
- [Kerberos Server Support, page 13-5](#)
- [LDAP Server Support, page 13-6](#)
- [SSO Support for WebVPN with HTTP Forms, page 13-9](#)
- [Local Database Support, page 13-9](#)

Summary of Support

[Table 13-1](#) summarizes the support for each AAA service by each AAA server type, including the local database. For more information about support for a specific AAA server type, refer to the topics following the table.

Table 13-1 Summary of AAA Support

AAA Service	Database Type							
	Local	RADIUS	TACACS+	SDI	NT	Kerberos	LDAP	HTTP Form
Authentication of...								
VPN users	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes ¹
Firewall sessions	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Administrators	Yes	Yes	Yes	Yes ²	Yes	Yes	Yes	No
Authorization of...								
VPN users	Yes	Yes	No	No	No	No	Yes	No
Firewall sessions	No	Yes ³	Yes	No	No	No	No	No
Administrators	Yes ⁴	No	Yes	No	No	No	No	No
Accounting of...								
VPN connections	No	Yes	Yes	No	No	No	No	No
Firewall sessions	No	Yes	Yes	No	No	No	No	No
Administrators	No	Yes ⁵	Yes	No	No	No	No	No

1. HTTP Form protocol supports single sign-on authentication for WebVPN users only.

2. SDI is not supported for HTTP administrative access.

3. For firewall sessions, RADIUS authorization is supported with user-specific access lists only, which are received or specified in a RADIUS authentication response.

4. Local command authorization is supported by privilege level only.

5. Command accounting is available for TACACS+ only.

RADIUS Server Support

The security appliance supports RADIUS servers.

This section contains the following topics:

- [Authentication Methods, page 13-4](#)
- [Attribute Support, page 13-4](#)
- [RADIUS Authorization Functions, page 13-4](#)

Authentication Methods

The security appliance supports the following authentication methods with RADIUS:

- PAP—For all connection types.
- CHAP—For L2TP-over-IPSec.
- MS-CHAPv1—For L2TP-over-IPSec.
- MS-CHAPv2—For L2TP-over-IPSec, and for regular IPSec remote access connections when the password management feature is enabled.

Attribute Support

The security appliance supports the following sets of RADIUS attributes:

- Authentication attributes defined in RFC 2138.
- Accounting attributes defined in RFC 2139.
- RADIUS attributes for tunneled protocol support, defined in RFC 2868.
- Cisco IOS VSAs, identified by RADIUS vendor ID 9.
- Cisco VPN-related VSAs, identified by RADIUS vendor ID 3076.
- Microsoft VSAs, defined in RFC 2548.

RADIUS Authorization Functions

The security appliance can use RADIUS servers for user authorization for network access using dynamic access lists or access list names per user. To implement dynamic access lists, you must configure the RADIUS server to support it. When the user authenticates, the RADIUS server sends a downloadable access list or access list name to the security appliance. Access to a given service is either permitted or denied by the access list. The security appliance deletes the access list when the authentication session expires.

TACACS+ Server Support

The security appliance supports TACACS+ authentication with ASCII, PAP, CHAP, and MS-CHAPv1.

SDI Server Support

The RSA SecureID servers are also known as SDI servers.

This section contains the following topics:

- [SDI Version Support, page 13-5](#)

- [Two-step Authentication Process, page 13-5](#)
- [SDI Primary and Replica Servers, page 13-5](#)

SDI Version Support

The security appliance supports SDI Version 5.0 and 6.0. SDI uses the concepts of an SDI primary and SDI replica servers. Each primary and its replicas share a single node secret file. The node secret file has its name based on the hexadecimal value of the ACE/Server IP address with .sdi appended.

A version 5.0 or 6.0 SDI server that you configure on the security appliance can be either the primary or any one of the replicas. See the [“SDI Primary and Replica Servers” section on page 13-5](#) for information about how the SDI agent selects servers to authenticate users.

Two-step Authentication Process

SDI version 5.0 and 6.0 uses a two-step process to prevent an intruder from capturing information from an RSA SecurID authentication request and using it to authenticate to another server. The Agent first sends a lock request to the SecurID server before sending the user authentication request. The server locks the username, preventing another (replica) server from accepting it. This means that the same user cannot authenticate to two security appliances using the same authentication servers simultaneously. After a successful username lock, the security appliance sends the passcode.

SDI Primary and Replica Servers

The security appliance obtains the server list when the first user authenticates to the configured server, which can be either a primary or a replica. The security appliance then assigns priorities to each of the servers on the list, and subsequent server selection derives at random from those assigned priorities. The highest priority servers have a higher likelihood of being selected.

NT Server Support

The security appliance supports Microsoft Windows server operating systems that support NTLM version 1, collectively referred to as NT servers.

**Note**

NT servers have a maximum length of 14 characters for user passwords. Longer passwords are truncated. This is a limitation of NTLM version 1.

Kerberos Server Support

The security appliance supports 3DES, DES, and RC4 encryption types.

**Note**

The security appliance does not support changing user passwords during tunnel negotiation. To avoid this situation happening inadvertently, disable password expiration on the Kerberos/Active Directory server for users connecting to the security appliance.

For a simple Kerberos server configuration example, see [Example 13-2](#).

LDAP Server Support

This section describes using an LDAP directory with the security appliance for user authentication and VPN authorization. This section includes the following topics:

- [Authentication with LDAP, page 13-6](#)
- [Authorization with LDAP for VPN, page 13-7](#)
- [LDAP Attribute Mapping, page 13-8](#)

For example configuration procedures used to set up LDAP authentication or authorization, see [Appendix E, “Configuring an External Server for Authorization and Authentication”](#).

Authentication with LDAP

During authentication, the security appliance acts as a client proxy to the LDAP server for the user, and authenticates to the LDAP server in either plain text or using the Simple Authentication and Security Layer (SASL) protocol. By default, the security appliance passes authentication parameters, usually a username and password, to the LDAP server in plain text. Whether using SASL or plain text, you can secure the communications between the security appliance and the LDAP server with SSL using the **ldap-over-ssl** command.

**Note**

If you do not configure SASL, we strongly recommend that you secure LDAP communications with SSL. See the **ldap-over-ssl** command in the *Cisco Security Appliance Command Reference*.

When user LDAP authentication has succeeded, the LDAP server returns the attributes for the authenticated user. For VPN authentication, these attributes generally include authorization data which is applied to the VPN session. Thus, using LDAP accomplishes authentication and authorization in a single step.

Securing LDAP Authentication with SASL

The security appliance supports the following SASL mechanisms, listed in order of increasing strength:

- Digest-MD5 — The security appliance responds to the LDAP server with an MD5 value computed from the username and password.
- Kerberos — The security appliance responds to the LDAP server by sending the username and realm using the GSSAPI (Generic Security Services Application Programming Interface) Kerberos mechanism.

You can configure the security appliance and LDAP server to support any combination of these SASL mechanisms. If you configure multiple mechanisms, the security appliance retrieves the list of SASL mechanisms configured on the server and sets the authentication mechanism to the strongest mechanism configured on both the security appliance and the server. For example, if both the LDAP server and the security appliance support both mechanisms, the security appliance selects Kerberos, the stronger of the mechanisms.

The following example configures the security appliance for authentication to an LDAP directory server named `ldap_dir_1` using the digest-MD5 SASL mechanism, and communicating over an SSL-secured connection:

```
hostname(config)# aaa-server ldap_dir_1 protocol ldap
hostname(config-aaa-server-group)# aaa-server ldap_dir_1 host 10.1.1.4
hostname(config-aaa-server-host)# sasl-mechanism digest-md5
hostname(config-aaa-server-host)# ldap-over-ssl enable
```

```
hostname(config-aaa-server-host) #
```

Setting the LDAP Server Type

The security appliance supports LDAP Version 3. In the current release, it is compatible only with the Sun Microsystems JAVA System Directory Server (formerly named the Sun ONE Directory Server) and the Microsoft Active Directory. In later releases, the security appliance will support other OpenLDAP servers.

By default, the security appliance auto-detects whether it is connected to a Microsoft or a Sun LDAP directory server. However, if auto-detection fails to determine the LDAP server type, and you know the server is either a Microsoft or Sun server, you can manually configure the server type. The following example sets the LDAP directory server `ldap_dir_1` to the Sun Microsystems type:

```
hostname(config) # aaa-server ldap_dir_1 protocol ldap
hostname(config-aaa-server-group) # aaa-server ldap_dir_1 host 10.1.1.4
hostname(config-aaa-server-host) # server-type sun
hostname(config-aaa-server-host) #
```



Note

- Sun—The DN configured on the security appliance to access a Sun directory server must be able to access the default password policy on that server. We recommend using the directory administrator, or a user with directory administrator privileges, as the DN. Alternatively, you can place an ACI on the default password policy.
- Microsoft—You must configure LDAP over SSL to enable password management with Microsoft Active Directory.

Authorization with LDAP for VPN

When user LDAP authentication for VPN access has succeeded, the security appliance queries the LDAP server which returns LDAP attributes. These attributes generally include authorization data that applies to the VPN session. Thus, using LDAP accomplishes authentication and authorization in a single step.

There may be cases, however, where you require authorization from an LDAP directory server that is separate and distinct from the authentication mechanism. For example, if you use an SDI or certificate server for authentication, no authorization information is passed back. For user authorizations in this case, you can query an LDAP directory after successful authentication, accomplishing authentication and authorization in two steps.

To set up VPN user authorization using LDAP, you must first create a AAA server group and a tunnel group. You then associate the server and tunnel groups using the **tunnel-group general-attributes** command. While there are other authorization-related commands and options available for specific requirements, the following example shows fundamental commands for enabling user authorization with LDAP. This example then creates an IPsec remote access tunnel group named `remote-1`, and assigns that new tunnel group to the previously created `ldap_dir_1` AAA server for authorization.

```
hostname(config) # tunnel-group remote-1 type ipsec-ra
hostname(config) # tunnel-group remote-1 general-attributes
hostname(config-general) # authorization-server-group ldap_dir_1
hostname(config-general) #
```

After you complete this fundamental configuration work, you can configure additional LDAP authorization parameters such as a directory password, a starting point for searching a directory, and the scope of a directory search:

```
hostname(config)# aaa-server ldap_dir_1 protocol ldap
hostname(config-aaa-server-group)# aaa-server ldap_dir_1 host 10.1.1.4
hostname(config-aaa-server-host)# ldap-login-dn obscurepassword
hostname(config-aaa-server-host)# ldap-base-dn starthere
hostname(config-aaa-server-host)# ldap-scope subtree
hostname(config-aaa-server-host)#
```

See LDAP commands in the *Cisco Security Appliance Command Reference* for more information.

LDAP Attribute Mapping

If you are introducing a security appliance to an existing LDAP directory, your existing LDAP attribute names and values are probably different from the existing ones. You must create LDAP attribute maps that map your existing user-defined attribute names and values to Cisco attribute names and values that are compatible with the security appliance. You can then bind these attribute maps to LDAP servers or remove them as needed. You can also show or clear attribute maps.



Note

To use the attribute mapping features correctly, you need to understand the Cisco LDAP attribute names and values as well as the user-defined attribute names and values.

The following command, entered in global configuration mode, creates an unpopulated LDAP attribute map table named att_map_1:

```
hostname(config)# ldap attribute-map att_map_1
hostname(config-ldap-attribute-map)#
```

The following commands map the user-defined attribute name department to the Cisco attribute name cVPN3000-IETF-Radius-Class. The second command maps the user-defined attribute value Engineering to the user-defined attribute department and the Cisco-defined attribute value group1.

```
hostname(config)# ldap attribute-map att_map_1
hostname(config-ldap-attribute-map)# map-name department cVPN3000-IETF-Radius-Class
hostname(config-ldap-attribute-map)# map-value department Engineering group1
hostname(config-ldap-attribute-map)#
```

The following commands bind the attribute map att_map_1 to the LDAP server ldap_dir_1:

```
hostname(config)# aaa-server ldap_dir_1 host 10.1.1.4
hostname(config-aaa-server-host)# ldap-attribute-map att_map_1
hostname(config-aaa-server-host)#
```



Note

The command to create an attribute map (**ldap attribute-map**) and the command to bind it to an LDAP server (**ldap-attribute-map**) differ only by a hyphen and the mode.

The following commands display or clear all LDAP attribute maps in the running configuration:

```
hostname# show running-config all ldap attribute-map
hostname(config)# clear configuration ldap attribute-map
hostname(config)#
```

The names of frequently mapped Cisco LDAP attributes and the type of user-defined attributes they would commonly be mapped to include:

```

cVPN3000-IETF-Radius-Class – Department or user group
cVPN3000-IETF-Radius-Filter-Id – Access control list
cVPN3000-IETF-Radius-Framed-IP-Address – A static IP address
cVPN3000-IPSec-Banner1 – A organization title
cVPN3000-Tunneling-Protocols – Allow or deny dial-in

```

For a list of Cisco LDAP attribute names and values, see [Appendix E, “Configuring an External Server for Authorization and Authentication”](#). Alternatively, you can enter “?” within ldap-attribute-map mode to display the complete list of Cisco LDAP attribute names, as shown in the following example:

```

hostname(config)# ldap attribute-map att_map_1
hostname(config-ldap-attribute-map)# map-name att_map_1 ?

ldap mode commands/options:
cisco-attribute-names:
  cVPN3000-Access-Hours
  cVPN3000-Allow-Network-Extension-Mode
  cVPN3000-Auth-Service-Type
  cVPN3000-Authenticated-User-Idle-Timeout
  cVPN3000-Authorization-Required
  cVPN3000-Authorization-Type
  :
  :
  cVPN3000-X509-Cert-Data
hostname(config-ldap-attribute-map)#

```

SSO Support for WebVPN with HTTP Forms

The security appliance can use the HTTP Form protocol for single sign-on (SSO) authentication of WebVPN users only. Single sign-on support lets WebVPN users enter a username and password only once to access multiple protected services and Web servers. The WebVPN server running on the security appliance acts as a proxy for the user to the authenticating server. When a user logs in, the WebVPN server sends an SSO authentication request, including username and password, to the authenticating server using HTTPS. If the server approves the authentication request, it returns an SSO authentication cookie to the WebVPN server. The security appliance keeps this cookie on behalf of the user and uses it to authenticate the user to secure websites within the domain protected by the SSO server.

In addition to the HTTP Form protocol, WebVPN administrators can choose to configure SSO with the HTTP Basic and NTLM authentication protocols (the **auto-signon** command), or with Computer Associates eTrust SiteMinder SSO server (formerly Netegrity SiteMinder) as well. For an in-depth discussion of configuring SSO with either HTTP Forms, **auto-signon** or SiteMinder, see the [Configuring WebVPN](#) chapter.

Local Database Support

The security appliance maintains a local database that you can populate with user profiles.

This section contains the following topics:

- [User Profiles, page 13-10](#)
- [Fallback Support, page 13-10](#)

User Profiles

User profiles contain, at a minimum, a username. Typically, a password is assigned to each username, although passwords are optional.

The **username attributes** command lets you enter the username mode. In this mode, you can add other information to a specific user profile. The information you can add includes VPN-related attributes, such as a VPN session timeout value.

Fallback Support

The local database can act as a fallback method for several functions. This behavior is designed to help you prevent accidental lockout from the security appliance.

For users who need fallback support, we recommend that their usernames and passwords in the local database match their usernames and passwords in the AAA servers. This provides transparent fallback support. Because the user cannot determine whether a AAA server or the local database is providing the service, using usernames and passwords on AAA servers that are different than the usernames and passwords in the local database means that the user cannot be certain which username and password should be given.

The local database supports the following fallback functions:

- **Console and enable password authentication**—When you use the **aaa authentication console** command, you can add the **LOCAL** keyword after the AAA server group tag. If the servers in the group all are unavailable, the security appliance uses the local database to authenticate administrative access. This can include enable password authentication, too.
- **Command authorization**—When you use the **aaa authorization command** command, you can add the **LOCAL** keyword after the AAA server group tag. If the TACACS+ servers in the group all are unavailable, the local database is used to authorize commands based on privilege levels.
- **VPN authentication and authorization**—VPN authentication and authorization are supported to enable remote access to the security appliance if AAA servers that normally support these VPN services are unavailable. The **authentication-server-group** command, available in tunnel-group general attributes mode, lets you specify the **LOCAL** keyword when you are configuring attributes of a tunnel group. When VPN client of an administrator specifies a tunnel group configured to fallback to the local database, the VPN tunnel can be established even if the AAA server group is unavailable, provided that the local database is configured with the necessary attributes.

Configuring the Local Database

This section describes how to manage users in the local database. You can use the local database for CLI access authentication, privileged mode authentication, command authorization, network access authentication, and VPN authentication and authorization. You cannot use the local database for network access authorization. The local database does not support accounting.

For multiple context mode, you can configure usernames in the system execution space to provide individual logins using the **login** command; however, you cannot configure any **aaa** commands in the system execution space.



Caution

If you add to the local database users who can gain access to the CLI but who should not be allowed to enter privileged mode, enable command authorization. (See the [“Configuring Local Command Authorization”](#) section on page 40-8.) Without command authorization, users can access privileged

mode (and all commands) at the CLI using their own password if their privilege level is 2 or greater (2 is the default). Alternatively, you can use RADIUS or TACACS+ authentication so that the user cannot use the **login** command, or you can set all local users to level 1 so you can control who can use the system enable password to access privileged mode.

To define a user account in the local database, perform the following steps:

Step 1 Create the user account. To do so, enter the following command:

```
hostname(config)# username name {nopassword | password password [mschap] } [privilege
priv_level]
```

where the options are as follows:

- **username**—A string from 4 to 64 characters long.
- **password** *password*—A string from 3 to 16 characters long.
- **mschap**—Specifies that the password will be converted to unicode and hashed using MD4 after you enter it. Use this keyword if users are authenticated using MSCHAPv1 or MSCHAPv2.
- **privilege level**—The privilege level that you want to assign to the new user account (from 0 to 15). The default is 2. This privilege level is used with command authorization.
- **nopassword**—Creates a user account with no password.

The **encrypted** and **nt-encrypted** keywords are typically for display only. When you define a password in the **username** command, the security appliance encrypts it when it saves it to the configuration for security purposes. When you enter the **show running-config** command, the **username** command does not show the actual password; it shows the encrypted password followed by the **encrypted** or **nt-encrypted** keyword (when you specify **mschap**). For example, if you enter the password “test,” the **show running-config** display would appear to be something like the following:

```
username pat password DLaUiAX3l78qgoB5c7iVNw== nt-encrypted
```

The only time you would actually enter the **encrypted** or **nt-encrypted** keyword at the CLI is if you are cutting and pasting a configuration to another security appliance and you are using the same password.

Step 2 To configure a local user account with VPN attributes, follow these steps:

a. Enter the following command:

```
hostname(config)# username username attributes
```

When you enter a **username attributes** command, you enter username mode. The commands available in this mode are as follows:

- **group-lock**
- **password-storage**
- **vpn-access-hours**
- **vpn-filter**
- **vpn-framed-ip-address**
- **vpn-group-policy**
- **vpn-idle-timeout**
- **vpn-session-timeout**
- **vpn-simultaneous-logins**
- **vpn-tunnel-protocol**

- **webvpn**

Use these commands as needed to configure the user profile. For more information about these commands, see the *Cisco Security Appliance Command Reference*.

- When you have finished configuring the user profiles, enter **exit** to return to config mode.

For example, the following command assigns a privilege level of 15 to the admin user account:

```
hostname(config)# username admin password passw0rd privilege 15
```

The following command creates a user account with no password:

```
hostname(config)# username bcham34 nopassword
```

The following commands create a user account with a password, enter username mode, and specify a few VPN attributes:

```
hostname(config)# username rwilliams password g0ge0us
hostname(config)# username rwilliams attributes
hostname(config-username)# vpn-tunnel-protocol IPSec
hostname(config-username)# vpn-simultaneous-logins 6
hostname(config-username)# exit
```

Identifying AAA Server Groups and Servers

If you want to use an external AAA server for authentication, authorization, or accounting, you must first create at least one AAA server group per AAA protocol and add one or more servers to each group. You identify AAA server groups by name. Each server group is specific to one type of server: Kerberos, LDAP, NT, RADIUS, SDI, or TACACS+.

The security appliance contacts the first server in the group. If that server is unavailable, the security appliance contacts the next server in the group, if configured. If all servers in the group are unavailable, the security appliance tries the local database if you configured it as a fallback method (management authentication and authorization only). If you do not have a fallback method, the security appliance continues to try the AAA servers.

To create a server group and add AAA servers to it, follow these steps:

Step 1 For each AAA server group you need to create, follow these steps:

- Identify the server group name and the protocol. To do so, enter the following command:

```
hostname(config)# aaa-server server_group protocol {kerberos | ldap | nt | radius |
sdi | tacacs+}
```

For example, to use RADIUS to authenticate network access and TACACS+ to authenticate CLI access, you need to create at least two server groups, one for RADIUS servers and one for TACACS+ servers.

You can have up to 15 single-mode server groups or 4 multi-mode server groups. Each server group can have up to 16 servers in single mode or up to 4 servers in multi-mode.

When you enter a **aaa-server protocol** command, you enter group mode.

- If you want to specify the maximum number of requests sent to a AAA server in the group before trying the next server, enter the following command:


```
hostname(config-aaa-server-group)# max-failed-attempts number
```

The *number* can be between 1 and 5. The default is 3.

If you configured a fallback method using the local database (for management access only; see the “Configuring AAA for System Administrators” section on page 40-5 and the “Configuring TACACS+ Command Authorization” section on page 40-11 to configure the fallback mechanism), and all the servers in the group fail to respond, then the group is considered to be unresponsive, and the fallback method is tried. The server group remains marked as unresponsive for a period of 10 minutes (by default) so that additional AAA requests within that period do not attempt to contact the server group, and the fallback method is used immediately. To change the unresponsive period from the default, see the **reactivation-mode** command in the following step.

If you do not have a fallback method, the security appliance continues to retry the servers in the group.

- c. If you want to specify the method (reactivation policy) by which failed servers in a group are reactivated, enter the following command:

```
hostname(config-aaa-server-group)# # reactivation-mode {depletion [deadtime minutes] | timed}
```

Where the **depletion** keyword reactivates failed servers only after all of the servers in the group are inactive.

The **deadtime** *minutes* argument specifies the amount of time in minutes, between 0 and 1440, that elapses between the disabling of the last server in the group and the subsequent re-enabling of all servers. The default is 10 minutes.

The **timed** keyword reactivates failed servers after 30 seconds of down time.

- d. If you want to send accounting messages to all servers in the group (RADIUS or TACACS+ only), enter the following command:

```
hostname(config-aaa-server-group)# accounting-mode simultaneous
```

To restore the default of sending messages only to the active server, enter the **accounting-mode single** command.

Step 2 For each AAA server on your network, follow these steps:

- a. Identify the server, including the AAA server group it belongs to. To do so, enter the following command:

```
hostname(config)# aaa-server server_group (interface_name) host server_ip
```

When you enter a **aaa-server host** command, you enter host mode.

- b. As needed, use host mode commands to further configure the AAA server.

The commands in host mode do not apply to all AAA server types. Table 13-2 lists the available commands, the server types they apply to, and whether a new AAA server definition has a default value for that command. Where a command is applicable to the server type you specified and no default value is provided (indicated by “—”), use the command to specify the value. For more information about these commands, see the *Cisco Security Appliance Command Reference*.

Table 13-2 *Host Mode Commands, Server Types, and Defaults*

Command	Applicable AAA Server Types	Default Value
accounting-port	RADIUS	1646
acl-netmask-convert	RADIUS	standard
authentication-port	RADIUS	1645
kerberos-realm	Kerberos	—
key	RADIUS	—
	TACACS+	—
ldap-attribute-map	LDAP	—
ldap-base-dn	LDAP	—
ldap-login-dn	LDAP	—
ldap-login-password	LDAP	—
ldap-naming-attribute	LDAP	—
ldap-over-ssl	LDAP	—
ldap-scope	LDAP	—
nt-auth-domain-controller	NT	—
radius-common-pw	RADIUS	—
retry-interval	Kerberos	10 seconds
	RADIUS	10 seconds
	SDI	10 seconds
sasl-mechanism	LDAP	—
server-port	Kerberos	88
	LDAP	389
	NT	139
	SDI	5500
	TACACS+	49
server-type	LDAP	auto-discovery
timeout	All	10 seconds

[Example 13-1](#) shows commands that add one TACACS+ group with one primary and one backup server, one RADIUS group with a single server, and an NT domain server.

Example 13-1 *Multiple AAA Server Groups and Servers*

```

hostname(config)# aaa-server AuthInbound protocol tacacs+
hostname(config-aaa-server-group)# max-failed-attempts 2
hostname(config-aaa-server-group)# reactivation-mode depletion deadtime 20
hostname(config-aaa-server-group)# exit
hostname(config)# aaa-server AuthInbound (inside) host 10.1.1.1
hostname(config-aaa-server-host)# key TACPlusUauthKey

```

```

hostname(config-aaa-server-host)# exit
hostname(config)# aaa-server AuthInbound (inside) host 10.1.1.2
hostname(config-aaa-server-host)# key TACPlusUauthKey2
hostname(config-aaa-server-host)# exit
hostname(config)# aaa-server AuthOutbound protocol radius
hostname(config-aaa-server-group)# exit
hostname(config)# aaa-server AuthOutbound (inside) host 10.1.1.3
hostname(config-aaa-server-host)# key RadUauthKey
hostname(config-aaa-server-host)# exit
hostname(config)# aaa-server NTAAuth protocol nt
hostname(config-aaa-server-group)# exit
hostname(config)# aaa-server NTAAuth (inside) host 10.1.1.4
hostname(config-aaa-server-host)# nt-auth-domain-controller primary1
hostname(config-aaa-server-host)# exit

```

Example 13-2 shows commands that configure a Kerberos AAA server group named `watchdogs`, add a AAA server to the group, and define the Kerberos realm for the server. Because **Example 13-2** does not define a retry interval or the port that the Kerberos server listens to, the security appliance uses the default values for these two server-specific parameters. **Table 13-2** lists the default values for all AAA server host mode commands.



Note

Kerberos realm names use numbers and upper-case letters only. Although the security appliance accepts lower-case letters for a realm name, it does not translate lower-case letters to upper-case letters. Be sure to use upper-case letters only.

Example 13-2 Kerberos Server Group and Server

```

hostname(config)# aaa-server watchdogs protocol kerberos
hostname(config-aaa-server-group)# aaa-server watchdogs host 192.168.3.4
hostname(config-aaa-server-host)# kerberos-realm EXAMPLE.COM
hostname(config-aaa-server-host)# exit
hostname(config)#

```

Using Certificates and User Login Credentials

The following section describes the different methods of using certificates and user login credentials (username and password) for authentication and authorization. This applies to both IPSec and WebVPN.

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.

Using User Login Credentials

The default method for authentication and authorization uses the user login credentials.

- Authentication
 - Enabled by authentication server group setting
 - Uses the username and password as credentials
- Authorization
 - Enabled by authorization server group setting
 - Uses the username as a credential

Using certificates

If user digital certificates are configured, the security appliance first validates the certificate. It does not, however, use any of the DN fields from the certificates as a username for the authentication.

If both authentication and authorization are enabled, the security appliance uses the user login credentials for both user authentication and authorization.

- Authentication
 - Enabled by authentication server group setting
 - Uses the username and password as credentials
- Authorization
 - Enabled by authorization server group setting
 - Uses the username as a credential

If authentication is disabled and authorization is enabled, the security appliance uses the primary DN field for authorization.

- Authentication
 - DISABLED (set to None) by authentication server group setting
 - No credentials used
- Authorization
 - Enabled by 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 security appliance uses the secondary DN field value as the username for the authorization request.

For example, consider a user certificate that contains 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.

Supporting a Zone Labs Integrity Server

This section introduces the Zone Labs Integrity Server, also called Check Point Integrity Server, and presents an example procedure for configuring the security appliance to support the Zone Labs Integrity Server. The Integrity server is a central management station for configuring and enforcing security policies on remote PCs. If a remote PC does not conform to the security policy dictated by the Integrity Server, it will not be granted access to the private network protected by the Integrity Server and security appliance.

This section includes the following topics:

- [Overview of Integrity Server and Security Appliance Interaction, page 13-17](#)
- [Configuring Integrity Server Support, page 13-17](#)

Overview of Integrity Server and Security Appliance Interaction

The VPN client software and the Integrity client software are co-resident on a remote PC. The following steps summarize the actions of the remote PC, security appliance, and Integrity server in the establishment of a session between the PC and the enterprise private network:

1. The VPN client software (residing on the same remote PC as the Integrity client software) connects to the security appliance and tells the security appliance what type of firewall client it is.
2. Once it approves the client firewall type, the security appliance passes Integrity server address information back to the Integrity client.
3. With the security appliance acting as a proxy, the Integrity client establishes a restricted connection with the Integrity server. A restricted connection is only between the Integrity client and server.
4. The Integrity server determines if the Integrity client is in compliance with the mandated security policies. If the client is in compliance with security policies, the Integrity server instructs the security appliance to open the connection and provide the client with connection details.
5. On the remote PC, the VPN client passes connection details to the Integrity client and signals that policy enforcement should begin immediately and the client can no enter the private network.
6. Once the connection is established, the server continues to monitor the state of the client using client heartbeat messages.

**Note**

The current release of the security appliance supports one Integrity Server at a time even though the user interfaces support the configuration of up to five Integrity Servers. If the active Server fails, configure another Integrity Server on the security appliance and then reestablish the client VPN session.

Configuring Integrity Server Support

This section describes an example procedure for configuring the security appliance to support the Zone Labs Integrity Servers. The procedure involves configuring address, port, connection fail timeout and fail states, and SSL certificate parameters.

First, you must configure the hostname or IP address of the Integrity server. The following example commands, entered in global configuration mode, configure an Integrity server using the IP address 10.0.0.5. They also specify port 300 (the default port is 5054) and the inside interface for communications with the Integrity server.

```
hostname(config)# zonelabs-integrity server-address 10.0.0.5
hostname(config)# zonelabs-integrity port 300
hostname(config)# zonelabs-integrity interface inside
hostname(config)#
```

If the connection between the security appliance and the Integrity server fails, the VPN client connections remain open by default so that the enterprise VPN is not disrupted by the failure of an Integrity server. However, you may want to close the VPN connections if the Zone Labs Integrity Server fails. The following commands ensure that the security appliance waits 12 seconds for a response from either the active or standby Integrity servers before declaring an the Integrity server as failed and closing the VPN client connections:

```
hostname(config)# zonelabs-integrity fail-timeout 12
hostname(config)# zonelabs-integrity fail-close
hostname(config)#
```

The following command returns the configured VPN client connection fail state to the default and ensures the client connections remain open:

```
hostname(config)# zonelabs-integrity fail-open  
hostname(config)#
```

The following example commands specify that the Integrity server connects to port 300 (default is port 80) on the security appliance to request the server SSL certificate. While the server SSL certificate is always authenticated, these commands also specify that the client SSL certificate of the Integrity server be authenticated.

```
hostname(config)# zonelabs-integrity ssl-certificate-port 300  
hostname(config)# zonelabs-integrity ssl-client-authentication  
hostname(config)#
```

To set the firewall client type to the Zone Labs Integrity type, use the **client-firewall** command as described in the [“Configuring Firewall Policies” section on page 30-55](#). The command arguments that specify firewall policies are not used when the firewall type is **zonelabs-integrity** because the Integrity server determines the policies.



CHAPTER 14

Configuring Failover

This chapter describes the security appliance failover feature, which lets you configure two security appliances so that one takes over operation if the other one fails.



Note

The ASA 5505 series adaptive security appliance does not support Stateful Failover or Active/Active failover.

This chapter includes the following sections:

- [Understanding Failover, page 14-1](#)
- [Configuring Failover, page 14-19](#)
- [Controlling and Monitoring Failover, page 14-49](#)

For failover configuration examples, see [Appendix B, “Sample Configurations.”](#)

Understanding Failover

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

The security appliance supports two failover configurations, Active/Active failover and Active/Standby failover. Each failover configuration has its own method for determining and performing failover.

With Active/Active failover, both units can pass network traffic. This lets you configure load balancing on your network. Active/Active failover is only available on units running in multiple context mode.

With Active/Standby failover, only one unit passes traffic while the other unit waits in a standby state. Active/Standby failover is available on units running in either single or multiple context mode.

Both failover configurations support stateful or stateless (regular) failover.



Note

VPN failover is not supported on units running in multiple context mode. VPN failover available for Active/Standby failover configurations only.

This section includes the following topics:

- [Failover System Requirements, page 14-2](#)
- [The Failover and Stateful Failover Links, page 14-3](#)
- [Active/Active and Active/Standby Failover, page 14-6](#)
- [Regular and Stateful Failover, page 14-15](#)
- [Failover Health Monitoring, page 14-16](#)
- [Failover Feature/Platform Matrix, page 14-18](#)
- [Failover Times by Platform, page 14-18](#)

Failover System Requirements

This section describes the hardware, software, and license requirements for security appliances in a failover configuration. This section contains the following topics:

- [Hardware Requirements, page 14-2](#)
- [Software Requirements, page 14-2](#)
- [License Requirements, page 14-2](#)

Hardware Requirements

The two units in a failover configuration must have the same hardware configuration. They must be the same model, have the same number and types of interfaces, and the same amount of RAM.



Note

The two units do not have to have the same size Flash memory. If 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 operating modes (routed or transparent, single or multiple context). They have the same major (first number) and minor (second number) software version. However, you can use different versions of the software during an upgrade process; for example, you can upgrade one unit from Version 7.0(1) to Version 7.0(2) and have failover remain active. We recommend upgrading both units to the same version to ensure long-term compatibility.

See [“Performing Zero Downtime Upgrades for Failover Pairs” section on page 41-6](#) for more information about upgrading the software on a failover pair.

License Requirements

On the PIX 500 series security appliance, at least one of the units must have an unrestricted (UR) license. The other unit can have a Failover Only (FO) license, a Failover Only Active-Active (FO_AA) license, or another UR license. Units with a Restricted license cannot be used for failover, and two units with FO or FO_AA licenses cannot be used together as a failover pair.

**Note**

The FO license does not support Active/Active failover.

The FO and FO_AA licenses are intended to be used solely for units in a failover configuration and not for units in standalone mode. If a failover unit with one of these licenses is used in standalone mode, the unit reboots at least once every 24 hours until the unit is returned to failover duty. A unit with an FO or FO_AA license operates in standalone mode if it is booted without being connected to a failover peer with a UR license. If the unit with a UR license in a failover pair fails and is removed from the configuration, the unit with the FO or FO_AA license does not automatically reboot every 24 hours; it operates uninterrupted unless the it is manually rebooted.

When the unit automatically reboots, the following message displays on the console:

```
=====NOTICE=====
      This machine is running in secondary mode without
      a connection to an active primary PIX. Please
      check your connection to the primary system.

                        REBOOTING....
=====
```

The ASA 5500 series adaptive security appliance platform does not have this restriction.

The Failover and Stateful Failover Links

This section describes the failover and the Stateful Failover links, which are dedicated connections between the two units in a failover configuration. This section includes the following topics:

- [Failover Link, page 14-3](#)
- [Stateful Failover Link, page 14-5](#)

Failover Link

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

- The unit state (active or standby).
- Power status (cable-based failover only—available only on the PIX 500 series security appliance).
- Hello messages (keep-alives).
- Network link status.
- MAC address exchange.
- Configuration replication and synchronization.

**Caution**

All information sent over the failover and Stateful Failover links is sent in clear text unless you secure the communication with a failover key. If the security appliance is used to terminate VPN tunnels, this information includes any usernames, passwords and preshared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communication with a failover key if you are using the security appliance to terminate VPN tunnels.

On the PIX 500 series security appliance, the failover link can be either a LAN-based connection or a dedicated serial Failover cable. On the ASA 5500 series adaptive security appliance, the failover link can only be a LAN-based connection.

This section includes the following topics:

- [LAN-Based Failover Link, page 14-4](#)
- [Serial Cable Failover Link \(PIX Security Appliance Only\), page 14-4](#)

LAN-Based Failover Link

You can use any unused Ethernet interface on the device as the failover link; however, you cannot specify an interface that is currently configured with a name. The LAN failover link interface is not configured as a normal networking interface. It exists for failover communication only. This interface should only be used for the LAN failover link (and optionally for the stateful failover link).

Connect the LAN 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 LAN failover interfaces of the ASA.
- Using a crossover Ethernet cable to connect the appliances directly, without the need for an external switch.



Note

When you use a crossover cable for the LAN failover link, if the LAN interface fails, the link is brought down on both peers. This condition may hamper troubleshooting efforts because you cannot easily determine which interface failed and caused the link to come down.



Note

The ASA supports Auto-MDI/MDIX on its copper Ethernet ports, so you can either use a crossover cable or a straight-through cable. If you use a straight-through cable, the interface automatically detects the cable and swaps one of the transmit/receive pairs to MDIX.

Serial Cable Failover Link (PIX Security Appliance Only)

The serial Failover cable, or “cable-based failover,” is only available on the PIX 500 series security appliance. If the two units are within six feet of each other, then we recommend that you use the serial Failover cable.

The cable that connects the two units is a modified RS-232 serial link cable that transfers data at 117,760 bps (115 Kbps). One end of the cable is labeled “Primary”. The unit attached to this end of the cable automatically becomes the primary unit. The other end of the cable is labeled “Secondary”. The unit attached to this end of the cable automatically becomes the secondary unit. You cannot override these designations in the PIX 500 series security appliance software. If you purchased a PIX 500 series security appliance failover bundle, this cable is included. To order a spare, use part number PIX-FO=.

The benefits of using cable-based failover include:

- The PIX 500 series security appliance can immediately detect a power loss on the peer unit and differentiate between a power loss from an unplugged cable.
- The standby unit can communicate with the active unit and can receive the entire configuration without having to be bootstrapped for failover. In LAN-based failover you need to configure the failover link on the standby unit before it can communicate with the active unit.
- The switch between the two units in LAN-based failover can be another point of hardware failure; cable-based failover eliminates this potential point of failure.

- You do not have to dedicate an Ethernet interface (and switch) to the failover link.
- The cable determines which unit is primary and which is secondary, eliminating the need to manually enter that information in the unit configurations.

The disadvantages include:

- Distance limitation—the units cannot be separated by more than 6 feet.
- Slower configuration replication.

Stateful Failover Link

To use Stateful Failover, you must configure a Stateful Failover link to pass all state information. You have three options for configuring a Stateful Failover link:

- You can use a dedicated Ethernet interface for the Stateful Failover link.
- If you are using LAN-based failover, you can share the failover link.
- You can share a regular data interface, such as the inside interface. However, this option is not recommended.

If you are using a dedicated Ethernet interface for the Stateful Failover link, you can use either a switch or a crossover cable to directly connect the units. If you use a switch, no other hosts or routers should be on this link.



Note

Enable the PortFast option on Cisco switch ports that connect directly to the security appliance.

If you use a data interface as the Stateful Failover link, you receive the following warning when you specify that interface as the Stateful Failover link:

```
***** WARNING ***** WARNING ***** WARNING ***** WARNING *****
Sharing Stateful failover interface with regular data interface is not
a recommended configuration due to performance and security concerns.
***** WARNING ***** WARNING ***** WARNING ***** WARNING *****
```

Sharing a data interface with the Stateful Failover interface can leave you vulnerable to replay attacks. Additionally, large amounts of Stateful Failover traffic may be sent on the interface, causing performance problems on that network segment.



Note

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

In multiple context mode, the Stateful Failover link resides in the system context. This interface and the failover interface are the only interfaces in the system context. All other interfaces are allocated to and configured from within security contexts.



Note

The IP address and MAC address for the Stateful Failover link does not change at failover unless the Stateful Failover link is configured on a regular data interface.



Caution

All information sent over the failover and Stateful Failover links is sent in clear text unless you secure the communication with a failover key. If the security appliance is used to terminate VPN tunnels, this information includes any usernames, passwords and preshared keys used for establishing the tunnels.

Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communication with a failover key if you are using the security appliance to terminate VPN tunnels.

Failover Interface Speed for Stateful Links

If you use the failover link as the Stateful Failover link, you should use the fastest Ethernet interface available. If you experience performance problems on that interface, consider dedicating a separate interface for the Stateful Failover interface.

Use the following failover interface speed guidelines for Cisco PIX security appliances and Cisco ASA adaptive security appliances:

- Cisco ASA 5520/5540/5550 and PIX 515E/535
 - The stateful link speed should match the fastest data link
- Cisco ASA 5510 and PIX 525
 - Stateful link speed can be 100 Mbps, even though the data interface can operate at 1 Gigabit due to the CPU speed limitation.

For optimum performance when using long distance LAN failover, the latency for the failover link should be less than 10 milliseconds and no more than 250 milliseconds. If latency is less than 10 milliseconds, some performance degradation occurs due to retransmission of failover messages.

All platforms support sharing of failover heartbeat and stateful link, but we recommend using a separate heartbeat link on systems with high Stateful Failover traffic.

Active/Active and Active/Standby Failover

This section describes each failover configuration in detail. This section includes the following topics:

- [Active/Standby Failover, page 14-6](#)
- [Active/Active Failover, page 14-10](#)
- [Determining Which Type of Failover to Use, page 14-15](#)

Active/Standby Failover

This section describes Active/Standby failover and includes the following topics:

- [Active/Standby Failover Overview, page 14-6](#)
- [Primary/Secondary Status and Active/Standby Status, page 14-7](#)
- [Device Initialization and Configuration Synchronization, page 14-7](#)
- [Command Replication, page 14-8](#)
- [Failover Triggers, page 14-9](#)
- [Failover Actions, page 14-9](#)

Active/Standby Failover Overview

Active/Standby failover lets you use a standby security appliance to take over the functionality of a failed unit. When the active unit fails, it changes to the standby state while the standby unit changes to the active state. The unit that becomes active assumes the IP addresses (or, for transparent firewall, the

management IP address) and MAC addresses of the failed unit and begins passing traffic. The unit that is now in standby state takes over the standby IP addresses and MAC addresses. Because network devices see no change in the MAC to IP address pairing, no ARP entries change or time out anywhere on the network.

**Note**

For multiple context mode, the security appliance can fail over the entire unit (including all contexts) but cannot fail over individual contexts separately.

Primary/Secondary Status and Active/Standby Status

The main differences between the two units in a failover pair are related to which unit is active and which unit is standby, namely which IP addresses to use and which unit actively passes traffic.

However, a few differences exist between the units based on which unit is primary (as specified in the configuration) and which unit is secondary:

- The primary unit always becomes the active unit if both units start up at the same time (and are of equal operational health).
- The primary unit MAC addresses are always coupled with the active IP addresses. The exception to this rule occurs when the secondary unit is active, and cannot obtain the primary unit MAC addresses over the failover link. In this case, the secondary unit MAC addresses are used.

Device Initialization and Configuration Synchronization

Configuration synchronization occurs when one or both devices in the failover pair boot. Configurations are always synchronized from the active unit to the standby unit. When the standby unit completes its initial startup, it clears its running configuration (except for the failover commands needed to communicate with the active unit), and the active unit sends its entire configuration to the standby unit.

The active unit is determined by the following:

- If a unit boots and detects a peer already running as active, it becomes the standby unit.
- If a unit boots and does not detect a peer, it becomes the active unit.
- If both units boot simultaneously, then the primary unit becomes the active unit and the secondary unit becomes the standby unit.

**Note**

If the secondary unit boots without detecting the primary unit, it becomes the active unit. It uses its own MAC addresses for the active IP addresses. However, when the primary unit becomes available, the secondary unit changes the MAC addresses to those of the primary unit, which can cause an interruption in your network traffic. To avoid this, configure the failover pair with virtual MAC addresses. See the [“Configuring Virtual MAC Addresses”](#) section on page 14-26 for more information.

When the replication starts, the security appliance console on the active unit displays the message “Beginning configuration replication: Sending to mate,” and when it is complete, the security appliance displays the message “End Configuration Replication to mate.” During replication, commands entered on the active unit may not replicate properly to the standby unit, and commands entered on the standby unit may be overwritten by the configuration being replicated from the active unit. Avoid entering commands on either unit in the failover pair during the configuration replication process. Depending upon the size of the configuration, replication can take from a few seconds to several minutes.

On the standby unit, the configuration exists only in running memory. To save the configuration to Flash memory after synchronization:

- For single context mode, enter the **write memory** command on the active unit. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory.
- For multiple context mode, enter the **write memory all** command on the active unit from the system execution space. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory. Using the **all** keyword with this command causes the system and all context configurations to be saved.

**Note**

Startup configurations saved on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts on disk from the active unit to an external server, and then copy them to disk on the standby unit, where they become available when the unit reloads.

Command Replication

Command replication always flows from the active unit to the standby unit. As commands are entered on the active unit, they are sent across the failover link to the standby unit. You do not have to save the active configuration to Flash memory to replicate the commands.

The following commands are replicated to the standby unit:

- all configuration commands except for the **mode**, **firewall**, and **failover lan unit** commands
- **copy running-config startup-config**
- **delete**
- **mkdir**
- **rename**
- **rmdir**
- **write memory**

The following commands are not replicated to the standby unit:

- 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**
- **mode**
- **show**

**Note**

Changes made on the standby unit are not replicated to the active unit. If you enter a command on the standby unit, the security appliance displays the message ****** WARNING **** Configuration Replication is NOT performed from Standby unit to Active unit. Configurations are no longer synchronized.** This message displays even when you enter many commands that do not affect the configuration.

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

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

Replicated commands are stored in the running configuration. To save the replicated commands to the Flash memory on the standby unit:

- For single context mode, enter the **copy running-config startup-config** command on the active unit. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory.
- For multiple context mode, enter the **copy running-config startup-config** command on the active unit from the system execution space and within each context on disk. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory. Contexts with startup configurations on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts on disk from the active unit to an external server, and then copy them to disk on the standby unit.

Failover Triggers

The unit can fail if one of the following events occurs:

- The unit has a hardware failure or a power failure.
- The unit has a software failure.
- Too many monitored interfaces fail.
- The **no failover active** command is entered on the active unit or the **failover active** command is entered on the standby unit.

Failover Actions

In Active/Standby failover, failover occurs on a unit basis. Even on systems running in multiple context mode, you cannot fail over individual or groups of contexts.

Table 14-1 shows the failover action for each failure event. For each failure event, the table shows the failover policy (failover or no failover), the action taken by the active unit, the action taken by the standby unit, and any special notes about the failover condition and actions.

Table 14-1 Failover Behavior

Failure Event	Policy	Active Action	Standby Action	Notes
Active unit failed (power or hardware)	Failover	n/a	Become active Mark active as failed	No hello messages are received on any monitored interface or the failover link.
Formerly active unit recovers	No failover	Become standby	No action	None.
Standby unit failed (power or hardware)	No failover	Mark standby as failed	n/a	When the standby unit is marked as failed, then the active unit does not attempt to fail over, even if the interface failure threshold is surpassed.
Failover link failed during operation	No failover	Mark failover interface as failed	Mark failover interface as failed	You should restore the failover link as soon as possible because the unit cannot fail over to the standby unit while the failover link is down.

Table 14-1 **Failover Behavior (continued)**

Failure Event	Policy	Active Action	Standby Action	Notes
Failover link failed at startup	No failover	Mark failover interface as failed	Become active	If the failover link is down at startup, both units become active.
Stateful Failover link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.
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.

Active/Active Failover

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

- [Active/Active Failover Overview, page 14-10](#)
- [Primary/Secondary Status and Active/Standby Status, page 14-11](#)
- [Device Initialization and Configuration Synchronization, page 14-11](#)
- [Command Replication, page 14-12](#)
- [Failover Triggers, page 14-13](#)
- [Failover Actions, page 14-14](#)

Active/Active Failover Overview

Active/Active failover is only available to security appliances in multiple context mode. In an Active/Active failover configuration, both security appliances can pass network traffic.

In Active/Active failover, you divide the security contexts on the security appliance into *failover groups*. A failover group is simply a logical group of one or more security contexts. You can create a maximum of two failover groups on the security appliance. The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default.

The failover group forms the base unit for failover in Active/Active failover. Interface failure monitoring, failover, and active/standby status are all attributes of a failover group rather than the unit. When an active failover group fails, it changes to the standby state while the standby failover group becomes active. The interfaces in the failover group that becomes active assume the MAC and IP addresses of the interfaces in the failover group that failed. The interfaces in the failover group that is now in the standby state take over the standby MAC and IP addresses.



Note

A failover group failing on a unit does not mean that the unit has failed. The unit may still have another failover group passing traffic on it.

When creating the failover groups, you should create them on the unit that will have failover group 1 in the active state.

**Note**

Active/Active failover generates virtual MAC addresses for the interfaces in each failover group. If you have more than one Active/Active failover pair on the same network, it is possible to have the same default virtual MAC addresses assigned to the interfaces on one pair as are assigned to the interfaces of the other pairs because of the way the default virtual MAC addresses are determined. To avoid having duplicate MAC addresses on your network, make sure you assign each physical interface a virtual active and standby MAC address.

Primary/Secondary Status and Active/Standby Status

As in Active/Standby failover, one unit in an Active/Active failover pair is designated the primary unit, and the other unit the secondary unit. Unlike Active/Standby failover, this designation does not indicate which unit becomes active when both units start simultaneously. Instead, the primary/secondary designation does two things:

- Determines which unit provides the running configuration to the pair when they boot simultaneously.
- Determines on which unit each failover group appears in the active state when the units boot simultaneously. Each failover group in the configuration is configured with a primary or secondary unit preference. You can configure both failover groups be in the active state on a single unit in the pair, with the other unit containing the failover groups in the standby state. However, a more typical configuration is to assign each failover group a different role preference to make each one active on a different unit, distributing the traffic across the devices.

**Note**

The security appliance does not provide load balancing services. Load balancing must be handled by a router passing traffic to the security appliance.

Which unit each failover group becomes active on is determined as follows:

- When a unit boots while the peer unit is not available, both failover groups become active on the unit.
- When a unit boots while the peer unit is active (with both failover groups in the active state), the failover groups remain in the active state on the active unit regardless of the primary or secondary preference of the failover group until one of the following:
 - A failover occurs.
 - You manually force the failover group to the other unit with the **no failover active** command.
 - You configured the failover group with the **preempt** command, which causes the failover group to automatically become active on the preferred unit when the unit becomes available.
- When both units boot at the same time, each failover group becomes active on its preferred unit after the configurations have been synchronized.

Device Initialization and Configuration Synchronization

Configuration synchronization occurs when one or both units in a failover pair boot. The configurations are synchronized as follows:

- When a unit boots while the peer unit is active (with both failover groups active on it), the booting unit contacts the active unit to obtain the running configuration regardless of the primary or secondary designation of the booting unit.

- When both units boot simultaneously, the secondary unit obtains the running configuration from the primary unit.

When the replication starts, the security appliance console on the unit sending the configuration displays the message “Beginning configuration replication: Sending to mate,” and when it is complete, the security appliance displays the message “End Configuration Replication to mate.” During replication, commands entered on the unit sending the configuration may not replicate properly to the peer unit, and commands entered on the unit receiving the configuration may be overwritten by the configuration being received. Avoid entering commands on either unit in the failover pair during the configuration replication process. Depending upon the size of the configuration, replication can take from a few seconds to several minutes.

On the unit receiving the configuration, the configuration exists only in running memory. To save the configuration to Flash memory after synchronization enter the **write memory all** command in the system execution space on the unit that has failover group 1 in the active state. The command is replicated to the peer unit, which proceeds to write its configuration to Flash memory. Using the **all** keyword with this command causes the system and all context configurations to be saved.

**Note**

Startup configurations saved on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts configuration files from the disk on the primary unit to an external server, and then copy them to disk on the secondary unit, where they become available when the unit reloads.

Command Replication

After both units are running, commands are replicated from one unit to the other as follows:

- Commands entered within a security context are replicated from the unit on which the security context appears in the active state to the peer unit.

**Note**

A context is considered in the active state on a unit if the failover group to which it belongs is in the active state on that unit.

- Commands entered in the system execution space are replicated from the unit on which failover group 1 is in the active state to the unit on which failover group 1 is in the standby state.
- Commands entered in the admin context are replicated from the unit on which failover group 1 is in the active state to the unit on which failover group 1 is in the standby state.

All configuration and file commands (**copy**, **rename**, **delete**, **mkdir**, **rmdir**, and so on) are replicated, with the following exceptions. The **show**, **debug**, **mode**, **firewall**, and **failover lan unit** commands are not replicated.

Failure to enter the commands on the appropriate unit for command replication to occur causes the configurations to be out of synchronization. Those changes may be lost the next time the initial configuration synchronization occurs.

The following commands are replicated to the standby unit:

- all configuration commands except for the **mode**, **firewall**, and **failover lan unit** commands
- **copy running-config startup-config**
- **delete**
- **mkdir**
- **rename**

- **rmdir**
- **write memory**

The following commands are not replicated to the standby unit:

- 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**
- **mode**
- **show**

You can use the **write standby** command to resynchronize configurations that have become out of sync. For Active/Active failover, the **write standby** command behaves as follows:

- If you enter the **write standby** command in the system execution space, the system configuration and the configurations for all of the security contexts on the security appliance is written to the peer unit. This includes configuration information for security contexts that are in the standby state. You must enter the command in the system execution space on the unit that has failover group 1 in the active state.

**Note**

If there are security contexts in the active state on the peer unit, the **write standby** command causes active connections through those contexts to be terminated. Use the **failover active** command on the unit providing the configuration to make sure all contexts are active on that unit before entering the **write standby** command.

- If you enter the **write standby** command in a security context, only the configuration for the security context is written to the peer unit. You must enter the command in the security context on the unit where the security context appears in the active state.

Replicated commands are not saved to the Flash memory when replicated to the peer unit. They are added to the running configuration. To save replicated commands to Flash memory on both units, use the **write memory** or **copy running-config startup-config** command on the unit that you made the changes on. The command is replicated to the peer unit and cause the configuration to be saved to Flash memory on the peer unit.

Failover Triggers

In Active/Active failover, failover can be triggered at the unit level if one of the following events occurs:

- The unit has a hardware failure.
- The unit has a power failure.
- The unit has a software failure.
- The **no failover active** or the **failover active** command is entered in the system execution space.

Failover is triggered at the failover group level when one of the following events occurs:

- Too many monitored interfaces in the group fail.
- The **no failover active group group_id** or **failover active group group_id** command is entered.

You configure the failover threshold for each failover group by specifying the number or percentage of interfaces within the failover group that must fail before the group fails. Because a failover group can contain multiple contexts, and each context can contain multiple interfaces, it is possible for all interfaces in a single context to fail without causing the associated failover group to fail.

See the “[Failover Health Monitoring](#)” section on page 14-16 for more information about interface and unit monitoring.

Failover Actions

In an Active/Active failover configuration, failover occurs on a failover group basis, not a system basis. For example, if you designate both failover groups as active on the primary unit, and failover group 1 fails, then failover group 2 remains active on the primary unit while failover group 1 becomes active on the secondary unit.



Note

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

Table 14-2 shows the failover action for each failure event. For each failure event, the policy (whether or not failover occurs), actions for the active failover group, and actions for the standby failover group are given.

Table 14-2 *Failover Behavior for Active/Active Failover*

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
A unit experiences a power or software failure	Failover	Become standby Mark as failed	Become active Mark active as failed	When a unit in a failover pair fails, any active failover groups on that unit are marked as failed and become active on the peer unit.
Interface failure on active failover group above threshold	Failover	Mark active group as failed	Become active	None.
Interface failure on standby failover group above threshold	No failover	No action	Mark standby group as failed	When the standby failover group is marked as failed, the active failover group does not attempt to fail over, even if the interface failure threshold is surpassed.
Formerly active failover group recovers	No failover	No action	No action	Unless configured with the preempt command, 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.

Table 14-2 Failover Behavior for Active/Active Failover (continued)

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
Stateful Failover link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.
Failover link failed during operation	No failover	n/a	n/a	Each unit marks the failover interface as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby unit while the failover link is down.

Determining Which Type of Failover to Use

The type of failover you choose depends upon your security appliance configuration and how you plan to use the security appliances.

If you are running the security appliance in single mode, then you can only use Active/Standby failover. Active/Active failover is only available to security appliances running in multiple context mode.

If you are running the security appliance in multiple context mode, then you can configure either Active/Active failover or Active/Standby failover.

- To provide load balancing, use Active/Active failover.
- If you do not want to provide load balancing, use Active/Standby or Active/Active failover.

Table 14-3 provides a comparison of some of the features supported by each type of failover configuration:

Table 14-3 Failover Configuration Feature Support

Feature	Active/Active	Active/Standby
Single Context Mode	No	Yes
Multiple Context Mode	Yes	Yes
Load Balancing Network Configurations	Yes	No
Unit Failover	Yes	Yes
Failover of Groups of Contexts	Yes	No
Failover of Individual Contexts	No	No

Regular and Stateful Failover

The security appliance supports two types of failover, regular and stateful. This section includes the following topics:

- [Regular Failover, page 14-16](#)
- [Stateful Failover, page 14-16](#)

Regular Failover

When a failover occurs, all active connections are dropped. Clients need to reestablish connections when the new active unit takes over.

Stateful Failover

When Stateful Failover is enabled, the active unit continually passes per-connection state information to the standby unit. After a failover occurs, the same connection information is available at the new active unit. Supported end-user applications are not required to reconnect to keep the same communication session.

The state information passed to the standby unit includes the following:

- NAT translation table.
- TCP connection states.
- UDP connection states.
- The ARP table.
- The Layer 2 bridge table (when running in transparent firewall mode).
- The HTTP connection states (if HTTP replication is enabled).
- The ISAKMP and IPsec SA table.
- GTP PDP connection database.

The information that is not passed to the standby unit when Stateful Failover is enabled includes the following:

- The HTTP connection table (unless HTTP replication is enabled).
- The user authentication (uauth) table.
- The routing tables. After a failover occurs, some packets may be lost or routed out of the wrong interface (the default route) while the dynamic routing protocols rediscover routes.
- State information for Security Service Modules.
- DHCP server address leases.
- L2TP over IPsec sessions.

**Note**

If failover occurs during an active Cisco IP SoftPhone session, the call remains active because the call session state information is replicated to the standby unit. When the call is terminated, the IP SoftPhone client loses connection with the Call Manager. This occurs because there is no session information for the CTIQBE hangup message on the standby unit. When the IP SoftPhone client does not receive a response back from the Call Manager within a certain time period, it considers the Call Manager unreachable and unregisters itself.

Failover Health Monitoring

The security appliance monitors each unit for overall health and for interface health. See the following sections for more information about how the security appliance performs tests to determine the state of each unit:

- [Unit Health Monitoring, page 14-17](#)

- [Interface Monitoring, page 14-17](#)

Unit Health Monitoring

The security appliance determines the health of the other unit by monitoring the failover link. When a unit does not receive three consecutive hello messages on the failover link, the unit sends an ARP request on all interfaces, including the failover interface. The action the security appliance takes depends on the response from the other unit. See the following possible actions:

- If the security appliance receives a response on the failover interface, then it does not fail over.
- If the security appliance does not receive a response on the failover link, but receives a response on another interface, then the unit does not failover. The failover link is marked as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby while the failover link is down.
- If the security appliance does not receive a response on any interface, then the standby unit switches to active mode and classifies the other unit as failed.



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.

You can configure the frequency of the hello messages and the hold time before failover occurs. A faster poll time and shorter hold time speed the detection of unit failures and make failover occur more quickly, but it can also cause “false” failures due to network congestion delaying the keepalive packets. See [Configuring Unit Health Monitoring, page 14-39](#) for more information about configuring unit health monitoring.

Interface Monitoring

You can monitor up to 250 interfaces divided between all contexts. You should monitor important interfaces, for example, you might configure one context to monitor a shared interface (because the interface is shared, all contexts benefit from the monitoring).

When a unit does not receive hello messages on a monitored interface for half of the configured hold time, it runs the following tests:

1. **Link Up/Down test**—A test of the interface status. If the Link Up/Down test indicates that the interface is operational, then the security appliance performs network tests. The purpose of these tests is to generate network traffic to determine which (if either) unit has failed. At the start of each test, each unit clears its received packet count for its interfaces. At the conclusion of each test, each unit looks to see if it has received any traffic. If it has, the interface is considered operational. If one unit receives traffic for a test and the other unit does not, the unit that received no traffic is considered failed. If neither unit has received traffic, then the next test is used.
2. **Network Activity test**—A received network activity test. The unit counts all received packets for up to 5 seconds. If any packets are received at any time during this interval, the interface is considered operational and testing stops. If no traffic is received, the ARP test begins.
3. **ARP test**—A reading of the unit ARP cache for the 2 most recently acquired entries. One at a time, the unit sends ARP requests to these machines, attempting to stimulate network traffic. After each request, the unit counts all received traffic for up to 5 seconds. If traffic is received, the interface is considered operational. If no traffic is received, an ARP request is sent to the next machine. If at the end of the list no traffic has been received, the ping test begins.

4. Broadcast Ping test—A ping test that consists of sending out a broadcast ping request. The unit then counts all received packets for up to 5 seconds. If any packets are received at any time during this interval, the interface is considered operational and testing stops.

If all network tests fail for an interface, but this interface on the other unit continues to successfully pass traffic, then the interface is considered to be failed. If the threshold for failed interfaces is met, then a failover occurs. If the other unit interface also fails all the network tests, then both interfaces go into the “Unknown” state and do not count towards the failover limit.

An interface becomes operational again if it receives any traffic. A failed security appliance returns to standby mode if the interface failure threshold is no longer met.

**Note**

If a failed unit does not recover and you believe it should not be failed, you can reset the state by entering the **failover reset** command. If the failover condition persists, however, the unit will fail again.

Failover Feature/Platform Matrix

Table 14-4 shows the failover features supported by each hardware platform.

Table 14-4 Failover Feature Support by Platform

Platform	Cable-Base Failover	LAN-Based Failover	Stateful Failover
ASA 5505 series adaptive security appliance	No	Yes	No
ASA 5500 series adaptive security appliance (other than the ASA 5505)	No	Yes	Yes
PIX 500 series security appliance	Yes	Yes	Yes

Failover Times by Platform

Table 14-5 shows the minimum, default, and maximum failover times for the PIX 500 series security appliance.

Table 14-5 PIX 500 series security appliance failover times.

Failover Condition	Minimum	Default	Maximum
Active unit loses power or stops normal operation.	800 milliseconds	45 seconds	45 seconds
Active unit interface link down.	500 milliseconds	5 seconds	15 seconds
Active unit interface up, but connection problem causes interface testing.	5 seconds	25 seconds	75 seconds

Table 14-6 shows the minimum, default, and maximum failover times for the ASA 5500 series adaptive security appliance.

Table 14-6 ASA 5500 series adaptive security appliance 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 card interface link down.	2 seconds	5 seconds	15 seconds
Active unit IPS or CSC card fails.	2 seconds	2 seconds	2 seconds
Active unit interface up, but connection problem causes interface testing.	5 seconds	25 seconds	75 seconds

Configuring Failover

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

- [Failover Configuration Limitations, page 14-19](#)
- [Configuring Active/Standby Failover, page 14-19](#)
- [Configuring Active/Active Failover, page 14-27](#)
- [Configuring Unit Health Monitoring, page 14-39](#)
- [Configuring Failover Communication Authentication/Encryption, page 14-39](#)
- [Verifying the Failover Configuration, page 14-40](#)

Failover Configuration Limitations

You cannot configure failover with the following type of IP addresses:

- IP addresses obtained through DHCP
- IP addresses obtained through PPPoE
- IPv6 addresses

Additionally, the following restrictions apply:

- Stateful Failover is not supported on the ASA 5505 adaptive security appliance.
- Active/Active failover is not supported on the ASA 5505 adaptive security appliance.
- You cannot configure failover when Easy VPN Remote is enabled on the ASA 5505 adaptive security appliance.
- VPN failover is not supported in multiple context mode.

Configuring Active/Standby Failover

This section provides step-by-step procedures for configuring Active/Standby failover. This section includes the following topics:

- [Prerequisites, page 14-20](#)
- [Configuring Cable-Based Active/Standby Failover \(PIX Security Appliance Only\), page 14-20](#)

- [Configuring LAN-Based Active/Standby Failover, page 14-21](#)
- [Configuring Optional Active/Standby Failover Settings, page 14-25](#)

Prerequisites

Before you begin, verify the following:

- Both units have the same hardware, software configuration, and proper license.
- Both units are in the same mode (single or multiple, transparent or routed).

Configuring Cable-Based Active/Standby Failover (PIX Security Appliance Only)

Follow these steps to configure Active/Standby failover using a serial cable as the failover link. The commands in this task are entered on the *primary* unit in the failover pair. The primary unit is the unit that has the end of the cable labeled “Primary” plugged into it. For devices in multiple context mode, the commands are entered in the system execution space unless otherwise noted.

You do not need to bootstrap the secondary unit in the failover pair when you use cable-based failover. Leave the secondary unit powered off until instructed to power it on.

Cable-based failover is only available on the PIX 500 series security appliance.

To configure cable-based Active/Standby failover, perform the following steps:

-
- Step 1** Connect the Failover cable to the PIX 500 series security appliances. Make sure that you attach the end of the cable marked “Primary” to the unit you use as the primary unit, and that you attach the end of the cable marked “Secondary” to the other unit.
- Step 2** Power on the primary unit.
- Step 3** If you have not done so already, configure the active and standby IP addresses for each data interface (routed mode), for the management IP address (transparent mode), or for the management-only interface. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces. The standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.



Note Do not configure an IP address for the Stateful Failover link if you are going to use a dedicated Stateful Failover interface. You use the **failover interface ip** command to configure a dedicated Stateful Failover interface in a later step.

```
hostname(config-if)# ip address active_addr netmask standby standby_addr
```

In routed firewall mode and for the management-only interface, this command is entered in interface configuration mode for each interface. In transparent firewall mode, the command is entered in global configuration mode.

In multiple context mode, you must configure the interface addresses from within each context. Use the **changeto context** command to switch between contexts. The command prompt changes to `hostname/context(config-if)#`, where *context* is the name of the current context. You must enter a management IP address for each context in transparent firewall multiple context mode.

- Step 4** (Optional) To enable Stateful Failover, configure the Stateful Failover link.



Note Stateful Failover is not available on the ASA 5505 series adaptive security appliance.

- a. Specify the interface to be used as the Stateful Failover link:

```
hostname(config)# failover link if_name phy_if
```

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface should not be used for any other purpose.

- b. Assign an active and standby IP address to the Stateful Failover link:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```



Note If the Stateful Failover link uses a data interface, skip this step. You have already defined the active and standby IP addresses for the interface.

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby IP address subnet mask.

The Stateful Failover link IP address and MAC address do not change at failover unless it uses a data interface. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

- c. Enable the interface:

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

Step 5 Enable failover:

```
hostname(config)# failover
```

Step 6 Power on the secondary unit and enable failover on the unit if it is not already enabled:

```
hostname(config)# failover
```

The active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages “Beginning configuration replication: sending to mate.” and “End Configuration Replication to mate” appear on the primary console.

Step 7 Save the configuration to Flash memory on the primary unit. Because the commands entered on the primary unit are replicated to the secondary unit, the secondary unit also saves its configuration to Flash memory.

```
hostname(config)# copy running-config startup-config
```

Configuring LAN-Based Active/Standby Failover

This section describes how to configure Active/Standby failover using an Ethernet failover link. When configuring LAN-based failover, you must bootstrap the secondary device to recognize the failover link before the secondary device can obtain the running configuration from the primary device.

**Note**

If you are changing from cable-based failover to LAN-based failover, you can skip any steps, such as assigning the active and standby IP addresses for each interface, that you completed for the cable-based failover configuration.

This section includes the following topics:

- [Configuring the Primary Unit, page 14-22](#)
- [Configuring the Secondary Unit, page 14-24](#)

Configuring the Primary Unit

Follow these steps to configure the primary unit in a LAN-based, Active/Standby failover configuration. These steps provide the minimum configuration needed to enable failover on the primary unit. For multiple context mode, all steps are performed in the system execution space unless otherwise noted.

To configure the primary unit in an Active/Standby failover pair, perform the following steps:

- Step 1** If you have not done so already, configure the active and standby IP addresses for each data interface (routed mode), for the management IP address (transparent mode), or for the management-only interface. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces. The standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.

**Note**

Do not configure an IP address for the Stateful Failover link if you are going to use a dedicated Stateful Failover interface. You use the **failover interface ip** command to configure a dedicated Stateful Failover interface in a later step.

```
hostname(config-if)# ip address active_addr netmask standby standby_addr
```

In routed firewall mode and for the management-only interface, this command is entered in interface configuration mode for each interface. In transparent firewall mode, the command is entered in global configuration mode.

In multiple context mode, you must configure the interface addresses from within each context. Use the **changeto context** command to switch between contexts. The command prompt changes to `hostname/context(config-if)#`, where *context* is the name of the current context. You must enter a management IP address for each context in transparent firewall multiple context mode.

- Step 2** (PIX security appliance only) Enable LAN-based failover:

```
hostname(config)# failover lan enable
```

- Step 3** Designate the unit as the primary unit:

```
hostname(config)# failover lan unit primary
```

- Step 4** Define the failover interface:

- a. Specify the interface to be used as the failover interface:

```
hostname(config)# failover lan interface if_name phy_if
```

The *if_name* argument assigns a name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. On the ASA 5505 adaptive security appliance, the *phy_if* specifies a VLAN.

- b. Assign the active and standby IP address to the failover link:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

The failover link IP address and MAC address do not change at failover. The active IP address for the failover link always stays with the primary unit, while the standby IP address stays with the secondary unit.

- c. Enable the interface:

```
hostname(config)# interface phy_if  
hostname(config-if)# no shutdown
```

Step 5 (Optional) To enable Stateful Failover, configure the Stateful Failover link.



Note Stateful Failover is not available on the ASA 5505 series adaptive security appliance.

- a. Specify the interface to be used as Stateful Failover link:

```
hostname(config)# failover link if_name phy_if
```



Note If the Stateful Failover link uses the failover link or a data interface, then you only need to supply the *if_name* argument.

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface should not be used for any other purpose (except, optionally, the failover link).

- b. Assign an active and standby IP address to the Stateful Failover link.



Note If the Stateful Failover link uses the failover link or data interface, skip this step. You have already defined the active and standby IP addresses for the interface.

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

The Stateful Failover link IP address and MAC address do not change at failover unless it uses a data interface. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

- c. Enable the interface.



Note If the Stateful Failover link uses the failover link or data interface, skip this step. You have already enabled the interface.

```
hostname(config)# interface phy_if  
hostname(config-if)# no shutdown
```

Step 6 Enable failover:

```
hostname(config)# failover
```

Step 7 Save the system configuration to Flash memory:

```
hostname(config)# copy running-config startup-config
```

Configuring the Secondary Unit

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

For multiple context mode, all steps are performed in the system execution space unless noted otherwise.

To configure the secondary unit, perform the following steps:

Step 1 (PIX security appliance only) Enable LAN-based failover:

```
hostname(config)# failover lan enable
```

Step 2 Define the failover interface. Use the same settings as you used for the primary unit.

a. Specify the interface to be used as the failover interface:

```
hostname(config)# failover lan interface if_name phy_if
```

The *if_name* argument assigns a name to the interface specified by the *phy_if* argument.

b. Assign the active and standby IP address to the failover link. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces.

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```



Note Enter this command exactly as you entered it on the primary unit when you configured the failover interface on the primary unit.

c. Enable the interface:

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

Step 3 (Optional) Designate this unit as the secondary unit:

```
hostname(config)# failover lan unit secondary
```



Note This step is optional because by default units are designated as secondary unless previously configured.

Step 4 Enable failover:

```
hostname(config)# failover
```

After you enable failover, the active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages “Beginning configuration replication: Sending to mate” and “End Configuration Replication to mate” appear on the active unit console.

Step 5 After the running configuration has completed replication, save the configuration to Flash memory:

```
hostname(config)# copy running-config startup-config
```

Configuring Optional Active/Standby Failover Settings

You can configure the following optional Active/Standby failover setting when you are initially configuring failover or after failover has already been configured. Unless otherwise noted, the commands should be entered on the active unit.

This section includes the following topics:

- [Enabling HTTP Replication with Stateful Failover, page 14-25](#)
- [Disabling and Enabling Interface Monitoring, page 14-25](#)
- [Configuring Interface Health Monitoring, page 14-26](#)
- [Configuring Failover Criteria, page 14-26](#)
- [Configuring Virtual MAC Addresses, page 14-26](#)

Enabling HTTP Replication with Stateful Failover

To allow HTTP connections to be included in the state information replication, you need to enable HTTP replication. Because HTTP connections are typically short-lived, and because HTTP clients typically retry failed connection attempts, HTTP connections are not automatically included in the replicated state information.

Enter the following command in global configuration mode to enable HTTP state replication when Stateful Failover is enabled:

```
hostname(config)# failover replication http
```

Disabling and Enabling Interface Monitoring

By default, monitoring physical interfaces is enabled and monitoring subinterfaces is disabled. You can monitor up to 250 interfaces on a unit. You can control which interfaces affect your failover policy by disabling the monitoring of specific interfaces and enabling the monitoring of others. This lets you exclude interfaces attached to less critical networks from affecting your failover policy.

For units in multiple configuration mode, use the following commands to enable or disable health monitoring for specific interfaces:

- To disable health monitoring for an interface, enter the following command within a context:

```
hostname/context(config)# no monitor-interface if_name
```

- To enable health monitoring for an interface, enter the following command within a context:

```
hostname/context(config)# monitor-interface if_name
```

For units in single configuration mode, use the following commands to enable or disable health monitoring for specific interfaces:

- To disable health monitoring for an interface, enter the following command in global configuration mode:

```
hostname(config)# no monitor-interface if_name
```

- To enable health monitoring for an interface, enter the following command in global configuration mode:

```
hostname(config)# monitor-interface if_name
```

Configuring Interface Health Monitoring

The security appliance sends hello packets out of each data interface to monitor interface health. If the security appliance does not receive a hello packet from the corresponding interface on the peer unit for over half of the hold time, then the additional interface testing begins. If a hello packet or a successful test result is not received within the specified hold time, the interface is marked as failed. Failover occurs if the number of failed interfaces meets the failover criteria.

Decreasing the poll and hold times enables the security appliance to detect and respond to interface failures more quickly, but may consume more system resources.

To change the interface poll time, enter the following command in global configuration mode:

```
hostname(config)# failover polltime interface [msec] time [holdtime time]
```

Valid values for the poll time are from 1 to 15 seconds or, if the optional msec keyword is used, from 500 to 999 milliseconds. The hold time determines how long it takes from the time a hello packet is missed to when the interface is marked as failed. Valid values for the hold time are from 5 to 75 seconds. You cannot enter a hold time that is less than 5 times the poll time.



Note

If the interface link is down, interface testing is not conducted and the standby unit could become active in just one interface polling period if the number of failed interface meets or exceeds the configured failover criteria.

Configuring Failover Criteria

By default, a single interface failure causes failover. You can specify a specific number of interfaces or a percentage of monitored interfaces that must fail before a failover occurs.

To change the default failover criteria, enter the following command in global configuration mode:

```
hostname(config)# failover interface-policy num[%]
```

When specifying a specific number of interfaces, the *num* argument can be from 1 to 250. When specifying a percentage of interfaces, the *num* argument can be from 1 to 100.

Configuring Virtual MAC Addresses

In Active/Standby failover, the MAC addresses for the primary unit are always associated with the active IP addresses. If the secondary unit boots first and becomes active, it uses the burned-in MAC address for its interfaces. When the primary unit comes online, the secondary unit obtains the MAC addresses from the primary unit. The change can disrupt network traffic.

You can configure virtual MAC addresses for each interface to ensure that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit. If you do not specify virtual MAC addresses the failover pair uses the burned-in NIC addresses as the MAC addresses.



Note

You cannot configure a virtual MAC address for the failover or Stateful Failover links. The MAC and IP addresses for those links do not change during failover.

Enter the following command on the active unit to configure the virtual MAC addresses for an interface:

```
hostname(config)# failover mac address phy_if active_mac standby_mac
```

The *phy_if* argument is the physical name of the interface, such as Ethernet1. The *active_mac* and *standby_mac* arguments are MAC addresses in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE.

The *active_mac* address is associated with the active IP address for the interface, and the *standby_mac* is associated with the standby IP address for the interface.

There are multiple ways to configure virtual MAC addresses on the security appliance. When more than one method has been used to configure virtual MAC addresses, the security appliance uses the following order of preference to determine which virtual MAC address is assigned to an interface:

1. The **mac-address** command (in interface configuration mode) address.
2. The **failover mac address** command address.
3. The **mac-address auto** command generated address.
4. The burned-in MAC address.

Use the **show interface** command to display the MAC address used by an interface.

Configuring Active/Active Failover

This section describes how to configure Active/Active failover.



Note

Active/Active failover is not available on the ASA 5505 series adaptive security appliance.

This section includes the following topics:

- [Prerequisites, page 14-27](#)
- [Configuring Cable-Based Active/Active Failover \(PIX security appliance\), page 14-27](#)
- [Configuring LAN-Based Active/Active Failover, page 14-29](#)
- [Configuring Optional Active/Active Failover Settings, page 14-33](#)

Prerequisites

Before you begin, verify the following:

- Both units have the same hardware, software configuration, and proper license.
- Both units are in multiple context mode.

Configuring Cable-Based Active/Active Failover (PIX security appliance)

Follow these steps to configure Active/Active failover using a serial cable as the failover link. The commands in this task are entered on the *primary* unit in the failover pair. The primary unit is the unit that has the end of the cable labeled “Primary” plugged into it. For devices in multiple context mode, the commands are entered in the system execution space unless otherwise noted.

You do not need to bootstrap the secondary unit in the failover pair when you use cable-based failover. Leave the secondary unit powered off until instructed to power it on.

Cable-based failover is only available on the PIX 500 series security appliance.

To configure cable-based, Active/Active failover, perform the following steps:

- Step 1** Connect the failover cable to the PIX 500 series security appliances. Make sure that you attach the end of the cable marked “Primary” to the unit you use as the primary unit, and that you attach the end of the cable marked “Secondary” to the unit you use as the secondary unit.
- Step 2** Power on the primary unit.
- Step 3** If you have not done so already, configure the active and standby IP addresses for each data interface (routed mode), for the management IP address (transparent mode), or for the management-only interface. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces. The standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.

You must configure the interface addresses from within each context. Use the **changeto context** command to switch between contexts. The command prompt changes to

`hostname/context(config-if)#`, where *context* is the name of the current context. You must enter a management IP address for each context in transparent firewall multiple context mode.



Note Do not configure an IP address for the Stateful Failover link if you are going to use a dedicated Stateful Failover interface. You use the **failover interface ip** command to configure a dedicated Stateful Failover interface in a later step.

```
hostname/context(config-if)# ip address active_addr netmask standby standby_addr
```

In routed firewall mode and for the management-only interface, this command is entered in interface configuration mode for each interface. In transparent firewall mode, the command is entered in global configuration mode.

- Step 4** (Optional) To enable Stateful Failover, configure the Stateful Failover link.

- a. Specify the interface to be used as Stateful Failover link:

```
hostname(config)# failover link if_name phy_if
```

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface should not be used for any other purpose (except, optionally, the failover link).

- b. Assign an active and standby IP address to the Stateful Failover link:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby IP address subnet mask.

The Stateful Failover link IP address and MAC address do not change at failover except for when Stateful Failover uses a regular data interface. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

- c. Enable the interface:

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

- Step 5** Configure the failover groups. You can have at most two failover groups. The **failover group** command creates the specified failover group if it does not exist and enters the failover group configuration mode.

For each failover group, you need to specify whether the failover group has primary or secondary preference using the **primary** or **secondary** command. You can assign the same preference to both failover groups. For load balancing configurations, you should assign each failover group a different unit preference.

The following example assigns failover group 1 a primary preference and failover group 2 a secondary preference:

```
hostname(config)# failover group 1
hostname(config-fover-group)# primary
hostname(config-fover-group)# exit
hostname(config)# failover group 2
hostname(config-fover-group)# secondary
hostname(config-fover-group)# exit
```

- Step 6** Assign each user context to a failover group using the **join-failover-group** command in context configuration mode.

Any unassigned contexts are automatically assigned to failover group 1. The admin context is always a member of failover group 1.

Enter the following commands to assign each context to a failover group:

```
hostname(config)# context context_name
hostname(config-context)# join-failover-group {1 | 2}
hostname(config-context)# exit
```

- Step 7** Enable failover:

```
hostname(config)# failover
```

- Step 8** Power on the secondary unit and enable failover on the unit if it is not already enabled:

```
hostname(config)# failover
```

The active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages “Beginning configuration replication: Sending to mate” and “End Configuration Replication to mate” appear on the primary console.

- Step 9** Save the configuration to Flash memory on the Primary unit. Because the commands entered on the primary unit are replicated to the secondary unit, the secondary unit also saves its configuration to Flash memory.

```
hostname(config)# copy running-config startup-config
```

- Step 10** If necessary, force any failover group that is active on the primary to the active state on the secondary. To force a failover group to become active on the secondary unit, issue the following command in the system execution space on the primary unit:

```
hostname# no failover active group group_id
```

The *group_id* argument specifies the group you want to become active on the secondary unit.

Configuring LAN-Based Active/Active Failover

This section describes how to configure Active/Active failover using an Ethernet failover link. When configuring LAN-based failover, you must bootstrap the secondary device to recognize the failover link before the secondary device can obtain the running configuration from the primary device.

This section includes the following topics:

- [Configure the Primary Unit, page 14-30](#)
- [Configure the Secondary Unit, page 14-32](#)

Configure the Primary Unit

To configure the primary unit in an Active/Active failover configuration, perform the following steps:

- Step 1** If you have not done so already, configure the active and standby IP addresses for each data interface (routed mode), for the management IP address (transparent mode), or for the management-only interface. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces. The standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.

You must configure the interface addresses from within each context. Use the **changeto context** command to switch between contexts. The command prompt changes to `hostname/context(config-if)#`, where *context* is the name of the current context. In transparent firewall mode, you must enter a management IP address for each context.



Note Do not configure an IP address for the Stateful Failover link if you are going to use a dedicated Stateful Failover interface. You use the **failover interface ip** command to configure a dedicated Stateful Failover interface in a later step.

```
hostname/context(config-if)# ip address active_addr netmask standby standby_addr
```

In routed firewall mode and for the management-only interface, this command is entered in interface configuration mode for each interface. In transparent firewall mode, the command is entered in global configuration mode.

- Step 2** Configure the basic failover parameters in the system execution space.

- a. (PIX security appliance only) Enable LAN-based failover:

```
hostname(config)# hostname(config)# failover lan enable
```

- b. Designate the unit as the primary unit:

```
hostname(config)# failover lan unit primary
```

- c. Specify the failover link:

```
hostname(config)# failover lan interface if_name phy_if
```

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. On the ASA 5505 adaptive security appliance, the *phy_if* specifies a VLAN. This interface should not be used for any other purpose (except, optionally, the Stateful Failover link).

- d. Specify the failover link active and standby IP addresses:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby IP address subnet mask. The failover link IP address and MAC address do not change at failover. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

Step 3 (Optional) To enable Stateful Failover, configure the Stateful Failover link:

- a. Specify the interface to be used as Stateful Failover link:

```
hostname(config)# failover link if_name phy_if
```

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface should not be used for any other purpose (except, optionally, the failover link).



Note If the Stateful Failover link uses the failover link or a regular data interface, then you only need to supply the *if_name* argument.

- b. Assign an active and standby IP address to the Stateful Failover link.



Note If the Stateful Failover link uses the failover link or a regular data interface, skip this step. You have already defined the active and standby IP addresses for the interface.

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

The state link IP address and MAC address do not change at failover. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

- c. Enable the interface.



Note If the Stateful Failover link uses the failover link or regular data interface, skip this step. You have already enabled the interface.

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

Step 4 Configure the failover groups. You can have at most two failover groups. The **failover group** command creates the specified failover group if it does not exist and enters the failover group configuration mode.

For each failover group, specify whether the failover group has primary or secondary preference using the **primary** or **secondary** command. You can assign the same preference to both failover groups. For load balancing configurations, you should assign each failover group a different unit preference.

The following example assigns failover group 1 a primary preference and failover group 2 a secondary preference:

```
hostname(config)# failover group 1
hostname(config-fover-group)# primary
hostname(config-fover-group)# exit
hostname(config)# failover group 2
hostname(config-fover-group)# secondary
hostname(config-fover-group)# exit
```

Step 5 Assign each user context to a failover group using the **join-failover-group** command in context configuration mode.

Any unassigned contexts are automatically assigned to failover group 1. The admin context is always a member of failover group 1.

Enter the following commands to assign each context to a failover group:

```
hostname(config)# context context_name
hostname(config-context)# join-failover-group {1 | 2}
hostname(config-context)# exit
```

Step 6 Enable failover:

```
hostname(config)# failover
```

Configure the Secondary Unit

When configuring LAN-based Active/Active failover, you need to bootstrap the secondary unit to recognize the failover link. This allows the secondary unit to communicate with and receive the running configuration from the primary unit.

To bootstrap the secondary unit in an Active/Active failover configuration, perform the following steps:

Step 1 (PIX security appliance only) Enable LAN-based failover:

```
hostname(config)# failover lan enable
```

Step 2 Define the failover interface. Use the same settings as you used for the primary unit:

- a. Specify the interface to be used as the failover interface:

```
hostname(config)# failover lan interface if_name phy_if
```

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. On the ASA 5505 adaptive security appliance, the *phy_if* specifies a VLAN.

- b. Assign the active and standby IP address to the failover link. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces.

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```



Note Enter this command exactly as you entered it on the primary unit when you configured the failover interface.

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

- c. Enable the interface:

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

Step 3 (Optional) Designate this unit as the secondary unit:

```
hostname(config)# failover lan unit secondary
```



Note This step is optional because by default units are designated as secondary unless previously configured otherwise.

Step 4 Enable failover:

```
hostname(config)# failover
```

After you enable failover, the active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages `Beginning configuration replication: Sending to mate` and `End Configuration Replication to mate` appear on the active unit console.

Step 5 After the running configuration has completed replication, enter the following command to save the configuration to Flash memory:

```
hostname(config)# copy running-config startup-config
```

Step 6 If necessary, force any failover group that is active on the primary to the active state on the secondary unit. To force a failover group to become active on the secondary unit, enter the following command in the system execution space on the primary unit:

```
hostname# no failover active group group_id
```

The *group_id* argument specifies the group you want to become active on the secondary unit.

Configuring Optional Active/Active Failover Settings

The following optional Active/Active failover settings can be configured when you are initially configuring failover or after you have already established failover. Unless otherwise noted, the commands should be entered on the unit that has failover group 1 in the active state.

This section includes the following topics:

- [Configuring Failover Group Preemption, page 14-33](#)
- [Enabling HTTP Replication with Stateful Failover, page 14-34](#)
- [Disabling and Enabling Interface Monitoring, page 14-34](#)
- [Configuring Interface Health Monitoring, page 14-34](#)
- [Configuring Failover Criteria, page 14-34](#)
- [Configuring Virtual MAC Addresses, page 14-35](#)
- [Configuring Asymmetric Routing Support, page 14-35](#)

Configuring Failover Group Preemption

Assigning a primary or secondary priority to a failover group specifies which unit the failover group becomes active on when both units boot simultaneously. However, if one unit boots before the other, then both failover groups become active on that unit. When the other unit comes online, any failover groups that have the unit as a priority do not become active on that unit unless manually forced over, a failover occurs, or the failover group is configured with the **preempt** command. The **preempt** command causes a failover group to become active on the designated unit automatically when that unit becomes available.

Enter the following commands to configure preemption for the specified failover group:

```
hostname(config)# failover group {1 | 2}  
hostname(config-fover-group)# preempt [delay]
```

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.

Enabling HTTP Replication with Stateful Failover

To allow HTTP connections to be included in the state information, you need to enable HTTP replication. Because HTTP connections are typically short-lived, and because HTTP clients typically retry failed connection attempts, HTTP connections are not automatically included in the replicated state information. You can use the **replication http** command to cause a failover group to replicate HTTP state information when Stateful Failover is enabled.

To enable HTTP state replication for a failover group, enter the following command. This command only affects the failover group in which it was configured. To enable HTTP state replication for both failover groups, you must enter this command in each group. This command should be entered in the system execution space.

```
hostname(config)# failover group {1 | 2}  
hostname(config-fover-group)# replication http
```

Disabling and Enabling Interface Monitoring

You can monitor up to 250 interfaces on a unit. By default, monitoring of physical interfaces is enabled and the monitoring of subinterfaces is disabled. You can control which interfaces affect your failover policy by disabling the monitoring of specific interfaces and enabling the monitoring of others. This lets you exclude interfaces attached to less critical networks from affecting your failover policy.

To disable health monitoring on an interface, enter the following command within a context:

```
hostname/context(config)# no monitor-interface if_name
```

To enable health monitoring on an interface, enter the following command within a context:

```
hostname/context(config)# monitor-interface if_name
```

Configuring Interface Health Monitoring

The security appliance sends hello packets out of each data interface to monitor interface health. If the security appliance does not receive a hello packet from the corresponding interface on the peer unit for over half of the hold time, then the additional interface testing begins. If a hello packet or a successful test result is not received within the specified hold time, the interface is marked as failed. Failover occurs if the number of failed interfaces meets the failover criteria.

Decreasing the poll and hold times enables the security appliance to detect and respond to interface failures more quickly, but may consume more system resources.

To change the default interface poll time, enter the following commands:

```
hostname(config)# failover group {1 | 2}  
hostname(config-fover-group)# polltime interface seconds
```

Valid values for the poll time are from 1 to 15 seconds or, if the optional msec keyword is used, from 500 to 999 milliseconds. The hold time determines how long it takes from the time a hello packet is missed to when the interface is marked as failed. Valid values for the hold time are from 5 to 75 seconds. You cannot enter a hold time that is less than 5 times the poll time.

Configuring Failover Criteria

By default, if a single interface fails failover occurs. You can specify a specific number of interfaces or a percentage of monitored interfaces that must fail before a failover occurs. The failover criteria is specified on a failover group basis.

To change the default failover criteria for the specified failover group, enter the following commands:

```
hostname(config)# failover group {1 | 2}
hostname(config-fover-group)# interface-policy num[%]
```

When specifying a specific number of interfaces, the *num* argument can be from 1 to 250. When specifying a percentage of interfaces, the *num* argument can be from 1 to 100.

Configuring Virtual MAC Addresses

Active/Active failover uses virtual MAC addresses on all interfaces. If you do not specify the virtual MAC addresses, then they are computed as follows:

- Active unit default MAC address: 00a0.c9*physical_port_number*.*failover_group_id*01.
- Standby unit default MAC address: 00a0.c9*physical_port_number*.*failover_group_id*02.



Note

If you have more than one Active/Active failover pair on the same network, it is possible to have the same default virtual MAC addresses assigned to the interfaces on one pair as are assigned to the interfaces of the other pairs because of the way the default virtual MAC addresses are determined. To avoid having duplicate MAC addresses on your network, make sure you assign each physical interface a virtual active and standby MAC address for all failover groups.

You can configure specific active and standby MAC addresses for an interface by entering the following commands:

```
hostname(config)# failover group {1 | 2}
hostname(config-fover-group)# mac address phy_if active_mac standby_mac
```

The *phy_if* argument is the physical name of the interface, such as Ethernet1. The *active_mac* and *standby_mac* arguments are MAC addresses in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE.

The *active_mac* address is associated with the active IP address for the interface, and the *standby_mac* is associated with the standby IP address for the interface.

There are multiple ways to configure virtual MAC addresses on the security appliance. When more than one method has been used to configure virtual MAC addresses, the security appliance uses the following order of preference to determine which virtual MAC address is assigned to an interface:

1. The **mac-address** command (in interface configuration mode) address.
2. The **failover mac address** command address.
3. The **mac-address auto** command generate address.
4. The automatically generated failover MAC address.

Use the **show interface** command to display the MAC address used by an interface.

Configuring Asymmetric Routing Support

When running in Active/Active failover, a unit may receive a return packet for a connection that originated through its peer unit. Because the security appliance that receives the packet does not have any connection information for the packet, the packet is dropped. This most commonly occurs when the two security appliances 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 using the **asr-group** command on interfaces where this is likely to occur. When an interface configured with the **asr-group** command receives a packet for which it has no session information, it checks the session information for the other interfaces that are in the same group. If it does not find a match, the packet is dropped. If it finds a match, then one of the following actions occurs:

- If the incoming traffic originated on a peer unit, 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**

Using the **asr-group** command to configure asymmetric routing support is more secure than using the **static** command with the **nailed** option.

The **asr-group** command does not provide asymmetric routing; it restores asymmetrically routed packets to the correct interface.

Prerequisites

You must have the following configured for asymmetric routing support to function properly:

- Active/Active Failover
- 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 security appliance to be able re-route asymmetrically routed HTTP packets, you need to replicate the HTTP state information.

You can configure the **asr-group** command on an interface without having failover configured, but it does not have any effect until Stateful Failover is enabled.

Configuring Support for Asymmetrically Routed Packets

To configure support for asymmetrically routed packets, perform the following steps:

Step 1 Configure Active/Active Stateful Failover for the failover pair. See [Configuring Active/Active Failover, page 14-27](#).

Step 2 For each interface that you want to participate in asymmetric routing support enter the following command. You must enter the command on the unit where the context is in the active state so that the command is replicated to the standby failover group. For more information about command replication, see [Command Replication, page 14-12](#).


```
hostname/ctx(config)# interface phy_if  
hostname/ctx(config-if)# asr-group num
```

Valid values for *num* range from 1 to 32. You need to enter the command for each interface that participates in the asymmetric routing group. You can view the number of ASR packets transmitted, received, or dropped by an interface using the **show interface detail** command. You can have more than one ASR group configured on the security appliance, but only one per interface. Only members of the same ASR group are checked for session information.

Example

Figure 14-1 shows an example of using the **asr-group** command for asymmetric routing support.

Figure 14-1 ASR Example



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

Example 14-1 Primary Unit System Configuration

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

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

Example 14-2 admin Context Configuration

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

Example 14-3 ctx1 Context Configuration

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

Figure 14-1 on page 14-37 shows the ASR support working as follows:

1. An outbound session passes through security appliance SecAppA. 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 security appliance SecAppB.
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 with the command **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 SecAppB. Stateful Failover replicated the session information from SecAppA to SecAppB.
5. Instead of being dropped, the layer 2 header is re-written 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.

Configuring Unit Health Monitoring

The security appliance sends hello packets over the failover interface to monitor unit health. If the standby unit does not receive a hello packet from the active unit for two consecutive polling periods, it sends additional testing packets through the remaining device interfaces. If a hello packet or a response to the interface test packets is not received within the specified hold time, the standby unit becomes active.

You can configure the frequency of hello messages when monitoring unit health. Decreasing the poll time allows a unit failure to be detected more quickly, but consumes more system resources.

To change the unit poll time, enter the following command in global configuration mode:

```
hostname(config)# failover polltime [msec] time [holdtime [msec] time]
```

You can configure the polling frequency from 1 to 15 seconds or, if the optional **msec** keyword is used, from 200 to 999 milliseconds. The hold time determines how long it takes from the time a hello packet is missed to when failover occurs. The hold time must be at least 3 times the poll time. You can configure the hold time from 1 to 45 seconds or, if the optional **msec** keyword is used, from 800 to 990 milliseconds.

Setting the security appliance to use the minimum poll and hold times allows it to detect and respond to unit failures in under a second, but it also increases system resource usage and can cause false failure detection in cases where the networks are congested or where the security appliance is running near full capacity.

Configuring Failover Communication Authentication/Encryption

You can encrypt and authenticate the communication between failover peers by specifying a shared secret or hexadecimal key.



Note

On the PIX 500 series security appliance, if you are using the dedicated serial failover cable to connect the units, then communication over the failover link is not encrypted even if a failover key is configured. The failover key only encrypts LAN-based failover communication.



Caution

All information sent over the failover and Stateful Failover links is sent in clear text unless you secure the communication with a failover key. If the security appliance is used to terminate VPN tunnels, this information includes any usernames, passwords and preshared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communication with a failover key if you are using the security appliance to terminate VPN tunnels.

Enter the following command on the active unit of an Active/Standby failover pair or on the unit that has failover group 1 in the active state of an Active/Active failover pair:

```
hostname(config)# failover key {secret | hex key}
```

The *secret* argument specifies a shared secret that is used to generate the encryption key. It can be from 1 to 63 characters. The characters can be any combination of numbers, letters, or punctuation. The **hex key** argument specifies a hexadecimal encryption key. The key must be 32 hexadecimal characters (0-9, a-f).

**Note**

To prevent the failover key from being replicated to the peer unit in clear text for an existing failover configuration, disable failover on the active unit (or in the system execution space on the unit that has failover group 1 in the active state), enter the failover key on both units, and then re-enable failover. When failover is re-enabled, the failover communication is encrypted with the key.

For new LAN-based failover configurations, the **failover key** command should be part of the failover pair bootstrap configuration.

Verifying the Failover Configuration

This section describes how to verify your failover configuration. This section includes the following topics:

- [Using the show failover Command, page 14-40](#)
- [Viewing Monitored Interfaces, page 14-48](#)
- [Displaying the Failover Commands in the Running Configuration, page 14-48](#)
- [Testing the Failover Functionality, page 14-49](#)

Using the show failover Command

This section describes the **show failover** command output. On each unit you can verify the failover status by entering the **show failover** command. The information displayed depends upon whether you are using Active/Standby or Active/Active failover.

This section includes the following topics:

- [show failover—Active/Standby, page 14-40](#)
- [Show Failover—Active/Active, page 14-44](#)

show failover—Active/Standby

The following is sample output from the **show failover** command for Active/Standby Failover. [Table 14-7](#) provides descriptions for the information shown.

```
hostname# show failover

Failover On
Cable status: N/A - LAN-based failover enabled
Failover unit Primary
Failover LAN Interface: fover Ethernet2 (up)
Unit Poll frequency 1 seconds, holdtime 3 seconds
Interface Poll frequency 15 seconds
Interface Policy 1
Monitored Interfaces 2 of 250 maximum
failover replication http
Last Failover at: 22:44:03 UTC Dec 8 2004
  This host: Primary - Active
    Active time: 13434 (sec)
    Interface inside (10.130.9.3): Normal
    Interface outside (10.132.9.3): Normal
  Other host: Secondary - Standby Ready
    Active time: 0 (sec)
    Interface inside (10.130.9.4): Normal
    Interface outside (10.132.9.4): Normal
```

```

Stateful Failover Logical Update Statistics
Link : fover Ethernet2 (up)
Stateful Obj    xmit    xerr    rcv    rerr
General        1950      0      1733      0
sys cmd        1733      0      1733      0
up time         0        0        0        0
RPC services    0        0        0        0
TCP conn        6        0        0        0
UDP conn        0        0        0        0
ARP tbl        106      0        0        0
Xlate_Timeout   0        0        0        0
VPN IKE upd     15        0        0        0
VPN IPSEC upd   90        0        0        0
VPN CTCP upd    0        0        0        0
VPN SDI upd     0        0        0        0
VPN DHCP upd    0        0        0        0

Logical Update Queue Information
          Cur    Max    Total
Recv Q:   0      2     1733
Xmit Q:   0      2    15225

```

In multiple context mode, using the **show failover** command in a security context displays the failover information for that context. The information is similar to the information shown when using the command in single context mode. Instead of showing the active/standby status of the unit, it displays the active/standby status of the context. [Table 14-7](#) provides descriptions for the information shown.

```

Failover On
Last Failover at: 04:03:11 UTC Jan 4 2003
  This context: Negotiation
    Active time: 1222 (sec)
    Interface outside (192.168.5.121): Normal
    Interface inside (192.168.0.1): Normal
  Peer context: Not Detected
    Active time: 0 (sec)
    Interface outside (192.168.5.131): Normal
    Interface inside (192.168.0.11): Normal

```

```

Stateful Failover Logical Update Statistics
Status: Configured.
Stateful Obj    xmit    xerr    rcv    rerr
RPC services    0        0        0        0
TCP conn        99        0        0        0
UDP conn        0        0        0        0
ARP tbl        22        0        0        0
Xlate_Timeout   0        0        0        0
GTP PDP         0        0        0        0
GTP PDMCB       0        0        0        0

```

Table 14-7 **Show Failover Display Description**

Field	Options
Failover	- On - Off
Cable status:	- Normal—The cable is connected to both units, and they both have power. - My side not connected—The serial cable is not connected to this unit. It is unknown if the cable is connected to the other unit. - Other side is not connected—The serial cable is connected to this unit, but not to the other unit. - Other side powered off—The other unit is turned off. - N/A—LAN-based failover is enabled.
Failover Unit	Primary or Secondary.
Failover LAN Interface	Displays the logical and physical name of the failover link.
Unit Poll frequency	Displays the number of seconds between hello messages sent to the peer unit and the number of seconds during which the unit must receive a hello message on the failover link before declaring the peer failed.
Interface Poll frequency	<i>n</i> seconds The number of seconds you set with the failover polltime interface command. The default is 15 seconds.
Interface Policy	Displays the number or percentage of interfaces that must fail to trigger failover.
Monitored Interfaces	Displays the number of interfaces monitored out of the maximum possible.
failover replication http	Displays if HTTP state replication is enabled for Stateful Failover.
Last Failover at:	The date and time of the last failover in the following form: <i>hh:mm:ss UTC DayName Month Day yyyy</i> UTC (Coordinated Universal Time) is equivalent to GMT (Greenwich Mean Time).
This host:	For each host, the display shows the following information.
Other host:	
Primary or Secondary	- Active - Standby
Active time:	<i>n</i> (sec) The amount of time the unit has been active. This time is cumulative, so the standby unit, if it was active in the past, also shows a value.
slot <i>x</i>	Information about the module in the slot or empty.

Table 14-7 Show Failover Display Description (continued)

Field	Options
Interface <i>name</i> (<i>n.n.n.n</i>):	For each interface, the display shows the IP address currently being used on each unit, as well as one of the following conditions: - Failed—The interface has failed. - No Link—The interface line protocol is down. - Normal—The interface is working correctly. - Link Down—The interface has been administratively shut down. - Unknown—The security appliance cannot determine the status of the interface. - Waiting—Monitoring of the network interface on the other unit has not yet started.
Stateful Failover Logical Update Statistics	The following fields relate to the Stateful Failover feature. If the Link field shows an interface name, the Stateful Failover statistics are shown.
Link	<i>interface_name</i>—The interface used for the Stateful Failover link. - Unconfigured—You are not using Stateful Failover. - up—The interface is up and functioning. - down—The interface is either administratively shutdown or is physically down. - failed—The interface has failed and is not passing stateful data.
Stateful Obj	For each field type, the following statistics are shown. They are counters for the number of state information packets sent between the two units; the fields do not necessarily show active connections through the unit. - xmit—Number of transmitted packets to the other unit. - xerr—Number of errors that occurred while transmitting packets to the other unit. - rcv—Number of received packets. - rerr—Number of errors that occurred while receiving packets from the other unit.
General	Sum of all stateful objects.
sys cmd	Logical update system commands; for example, LOGIN and Stay Alive.
up time	Up time, which the active unit passes to the standby unit.
RPC services	Remote Procedure Call connection information.
TCP conn	TCP connection information.
UDP conn	Dynamic UDP connection information.
ARP tbl	Dynamic ARP table information.
L2BRIDGE tbl	Layer 2 bridge table information (transparent firewall mode only).
Xlate_Timeout	Indicates connection translation timeout information.
VPN IKE upd	IKE connection information.

Table 14-7 *Show Failover Display Description (continued)*

Field	Options
VPN IPSEC upd	IPSec connection information.
VPN CTCP upd	cTCP tunnel connection information.
VPN SDI upd	SDI AAA connection information.
VPN DHCP upd	Tunneled DHCP connection information.
GTP PDP	GTP PDP update information. This information appears only if inspect GTP is enabled.
GTP PDPMCB	GTP PDPMCB update information. This information appears only if inspect GTP is enabled.
Logical Update Queue Information	For each field type, the following statistics are used: • Cur—Current number of packets • Max—Maximum number of packets • Total—Total number of packets
Recv Q	The status of the receive queue.
Xmit Q	The status of the transmit queue.

Show Failover—Active/Active

The following is sample output from the **show failover** command for Active/Active Failover. [Table 14-8](#) provides descriptions for the information shown.

```
hostname# show failover

Failover On
Failover unit Primary
Failover LAN Interface: third GigabitEthernet0/2 (up)
Unit Poll frequency 1 seconds, holdtime 15 seconds
Interface Poll frequency 4 seconds
Interface Policy 1
Monitored Interfaces 8 of 250 maximum
failover replication http
Group 1 last failover at: 13:40:18 UTC Dec 9 2004
Group 2 last failover at: 13:40:06 UTC Dec 9 2004

This host:      Primary
Group 1         State:      Active
                Active time: 2896 (sec)
Group 2         State:      Standby Ready
                Active time:  0 (sec)

slot 0: ASA-5530 hw/sw rev (1.0/7.0(0)79) status (Up Sys)
slot 1: SSM-IDS-20 hw/sw rev (1.0/5.0(0.11)S91(0.11)) status (Up)
admin Interface outside (10.132.8.5): Normal
admin Interface third (10.132.9.5): Normal
admin Interface inside (10.130.8.5): Normal
admin Interface fourth (10.130.9.5): Normal
ctx1 Interface outside (10.1.1.1): Normal
ctx1 Interface inside (10.2.2.1): Normal
ctx2 Interface outside (10.3.3.2): Normal
ctx2 Interface inside (10.4.4.2): Normal

Other host:     Secondary
```

```

Group 1      State:      Standby Ready
             Active time: 190 (sec)
Group 2      State:      Active
             Active time: 3322 (sec)

slot 0: ASA-5530 hw/sw rev (1.0/7.0(0)79) status (Up Sys)
slot 1: SSM-IDS-20 hw/sw rev (1.0/5.0(0.1)S91(0.1)) status (Up)
admin Interface outside (10.132.8.6): Normal
admin Interface third (10.132.9.6): Normal
admin Interface inside (10.130.8.6): Normal
admin Interface fourth (10.130.9.6): Normal
ctx1 Interface outside (10.1.1.2): Normal
ctx1 Interface inside (10.2.2.2): Normal
ctx2 Interface outside (10.3.3.1): Normal
ctx2 Interface inside (10.4.4.1): Normal

```

Stateful Failover Logical Update Statistics

```

Link : third GigabitEthernet0/2 (up)
Stateful Obj  xmit      xerr      rcv      rerr
General       1973       0       1895     0
sys cmd       380        0       380      0
up time       0          0        0        0
RPC services  0           0        0        0
TCP conn      1435       0       1450     0
UDP conn      0          0        0        0
ARP tbl       124        0        65       0
Xlate_Timeout 0           0        0        0
VPN IKE upd   15          0        0        0
VPN IPSEC upd 90          0        0        0
VPN CTCP upd  0           0        0        0
VPN SDI upd   0           0        0        0
VPN DHCP upd  0           0        0        0

```

Logical Update Queue Information

```

          Cur      Max      Total
Recv Q:   0        1       1895
Xmit Q:   0        0       1940

```

The following is sample output from the **show failover group** command for Active/Active Failover. The information displayed is similar to that of the **show failover** command, but limited to the specified group. [Table 14-8](#) provides descriptions for the information shown.

hostname# **show failover group 1**

Last Failover at: 04:09:59 UTC Jan 4 2005

```

This host:   Secondary
             State:      Active
             Active time: 186 (sec)

             admin Interface outside (192.168.5.121): Normal
             admin Interface inside (192.168.0.1): Normal

```

```

Other host:  Primary
             State:      Standby
             Active time: 0 (sec)

             admin Interface outside (192.168.5.131): Normal
             admin Interface inside (192.168.0.11): Normal

```

```

Stateful Failover Logical Update Statistics
Status: Configured.

```

```

RPC services      0          0          0          0
TCP conn          33         0          0          0
UDP conn          0          0          0          0
ARP tbl           12         0          0          0
Xlate_Timeout     0          0          0          0
GTP PDP           0          0          0          0
GTP PDPMCB        0          0          0          0

```

Table 14-8 *Show Failover Display Description*

Field	Options
Failover	- On - Off
Failover Unit	Primary or Secondary.
Failover LAN Interface	Displays the logical and physical name of the failover link.
Unit Poll frequency	Displays the number of seconds between hello messages sent to the peer unit and the number of seconds during which the unit must receive a hello message on the failover link before declaring the peer failed.
Interface Poll frequency	<i>n</i> seconds The number of seconds you set with the failover polltime interface command. The default is 15 seconds.
Interface Policy	Displays the number or percentage of interfaces that must fail before triggering failover.
Monitored Interfaces	Displays the number of interfaces monitored out of the maximum possible.
Group 1 Last Failover at: Group 2 Last Failover at:	The date and time of the last failover for each group in the following form: <i>hh:mm:ss UTC DayName Month Day yyyy</i> UTC (Coordinated Universal Time) is equivalent to GMT (Greenwich Mean Time).
This host: Other host:	For each host, the display shows the following information.
Role	Primary or Secondary
System State	- Active or Standby Ready - Active Time in seconds
Group 1 State Group 2 State	- Active or Standby Ready - Active Time in seconds
slot <i>x</i>	Information about the module in the slot or empty.

Table 14-8 Show Failover Display Description (continued)

Field	Options
<i>context</i> Interface name (<i>n.n.n.n</i>):	For each interface, the display shows the IP address currently being used on each unit, as well as one of the following conditions: - Failed—The interface has failed. - No link—The interface line protocol is down. - Normal—The interface is working correctly. - Link Down—The interface has been administratively shut down. - Unknown—The security appliance cannot determine the status of the interface. - Waiting—Monitoring of the network interface on the other unit has not yet started.
Stateful Failover Logical Update Statistics	The following fields relate to the Stateful Failover feature. If the Link field shows an interface name, the Stateful Failover statistics are shown.
Link	<i>interface_name</i>—The interface used for the Stateful Failover link. - Unconfigured—You are not using Stateful Failover. - up—The interface is up and functioning. - down—The interface is either administratively shutdown or is physically down. - failed—The interface has failed and is not passing stateful data.
Stateful Obj	For each field type, the following statistics are used. They are counters for the number of state information packets sent between the two units; the fields do not necessarily show active connections through the unit. - xmit—Number of transmitted packets to the other unit - xerr—Number of errors that occurred while transmitting packets to the other unit - rcv—Number of received packets - rerr—Number of errors that occurred while receiving packets from the other unit
General	Sum of all stateful objects.
sys cmd	Logical update system commands; for example, LOGIN and Stay Alive.
up time	Up time, which the active unit passes to the standby unit.
RPC services	Remote Procedure Call connection information.
TCP conn	TCP connection information.
UDP conn	Dynamic UDP connection information.
ARP tbl	Dynamic ARP table information.
L2BRIDGE tbl	Layer 2 bridge table information (transparent firewall mode only).
Xlate_Timeout	Indicates connection translation timeout information.
VPN IKE upd	IKE connection information.

Table 14-8 *Show Failover Display Description (continued)*

Field	Options
VPN IPSEC upd	IPSec connection information.
VPN CTCP upd	cTCP tunnel connection information.
VPN SDI upd	SDI AAA connection information.
VPN DHCP upd	Tunneled DHCP connection information.
GTP PDP	GTP PDP update information. This information appears only if inspect GTP is enabled.
GTP PDPMCB	GTP PDPMCB update information. This information appears only if inspect GTP is enabled.
Logical Update Queue Information	For each field type, the following statistics are used: • Cur—Current number of packets • Max—Maximum number of packets • Total—Total number of packets
Recv Q	The status of the receive queue.
Xmit Q	The status of the transmit queue.

Viewing Monitored Interfaces

To view the status of monitored interfaces, enter the following command. In single context mode, enter this command in global configuration mode. In multiple context mode, enter this command within a context.

```
primary/context(config)# show monitor-interface
```

For example:

```
hostname/context(config)# show monitor-interface
This host: Primary - Active
    Interface outside (192.168.1.2): Normal
    Interface inside (10.1.1.91): Normal
Other host: Secondary - Standby
    Interface outside (192.168.1.3): Normal
    Interface inside (10.1.1.100): Normal
```

Displaying the Failover Commands in the Running Configuration

To view the failover commands in the running configuration, enter the following command:

```
hostname(config)# show running-config failover
```

All of the failover commands are displayed. On units running multiple context mode, enter this command in the system execution space. Entering **show running-config all failover** displays the failover commands in the running configuration and includes commands for which you have not changed the default value.

Testing the Failover Functionality

To test failover functionality, perform the following steps:

-
- Step 1** Test that your active unit or failover group is passing traffic as expected by using FTP (for example) to send a file between hosts on different interfaces.
- Step 2** Force a failover to the standby unit by entering the following command:
- For Active/Standby failover, enter the following command on the active unit:
`hostname(config)# no failover active`
 - For Active/Active failover, enter the following command on the unit where the failover group containing the interface connecting your hosts is active:
`hostname(config)# no failover active group group_id`
- Step 3** Use FTP to send another file between the same two hosts.
- Step 4** If the test was not successful, enter the **show failover** command to check the failover status.
- Step 5** When you are finished, you can restore the unit or failover group to active status by enter the following command:
- For Active/Standby failover, enter the following command on the active unit:
`hostname(config)# failover active`
 - For Active/Active failover, enter the following command on the unit where the failover group containing the interface connecting your hosts is active:
`hostname(config)# failover active group group_id`
-

Controlling and Monitoring Failover

This sections describes how to control and monitor failover. This section includes the following topics:

- [Forcing Failover, page 14-49](#)
- [Disabling Failover, page 14-50](#)
- [Restoring a Failed Unit or Failover Group, page 14-50](#)
- [Monitoring Failover, page 14-50](#)

Forcing Failover

To force the standby unit or failover group to become active, enter one of the following commands:

- For Active/Standby failover:
Enter the following command on the standby unit:

```
hostname# failover active
```

Or, enter the following command on the active unit:

```
hostname# no failover active
```

- For Active/Active failover:

Enter the following command in the system execution space of the unit where the failover group is in the standby state:

```
hostname# failover active group group_id
```

Or, enter the following command in the system execution space of the unit where the failover group is in the active state:

```
hostname# no failover active group group_id
```

Entering the following command in the system execution space causes all failover groups to become active:

```
hostname# failover active
```

Disabling Failover

To disable failover, enter the following command:

```
hostname(config)# no failover
```

Disabling failover on an Active/Standby pair causes the active and standby state of each unit to be maintained until you restart. For example, the standby unit remains in standby mode so that both units do not start passing traffic. To make the standby unit active (even with failover disabled), see the [“Forcing Failover” section on page 14-49](#).

Disabling failover on an Active/Active pair causes the failover groups to remain in the active state on whichever unit they are currently active on, no matter which unit they are configured to prefer. The no failover command should be entered in the system execution space.

Restoring a Failed Unit or Failover Group

To restore a failed unit to an unfailed state, enter the following command:

```
hostname(config)# failover reset
```

To restore a failed Active/Active failover group to an unfailed state, enter the following command:

```
hostname(config)# failover reset group group_id
```

Restoring a failed unit or group to an unfailed state does not automatically make it active; restored units or groups remain in the standby state until made active by failover (forced or natural). An exception is a failover group configured with the **preempt** command. If previously active, a failover group becomes active if it is configured with the **preempt** command and if the unit on which it failed is the preferred unit.

Monitoring Failover

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

- [Failover System Messages, page 14-51](#)

- [Debug Messages, page 14-51](#)
- [SNMP, page 14-51](#)

Failover System Messages

The security appliance issues a number of system messages related to failover at priority level 2, which indicates a critical condition. To view these messages, see the *Cisco Security Appliance Logging Configuration and System Log Messages* to enable logging and to see descriptions of the system messages.



Note

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

Debug Messages

To see debug messages, enter the **debug fover** command. See the *Cisco Security Appliance Command Reference* for more information.



Note

Because debugging output is assigned high priority in the CPU process, it can drastically affect system performance. For this reason, use the **debug fover** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco TAC.

SNMP

To receive SNMP syslog traps for failover, configure the SNMP agent to send SNMP traps to SNMP management stations, define a syslog host, and compile the Cisco syslog MIB into your SNMP management station. See the **snmp-server** and **logging** commands in the *Cisco Security Appliance Command Reference* for more information.



PART 2

Configuring the Firewall



CHAPTER 15

Firewall Mode Overview

This chapter describes how the firewall works in each firewall mode. To set the firewall mode, see the [“Setting Transparent or Routed Firewall Mode”](#) section on page 2-5.



Note

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

This chapter includes the following sections:

- [Routed Mode Overview, page 15-1](#)
- [Transparent Mode Overview, page 15-8](#)

Routed Mode Overview

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

This section includes the following topics:

- [IP Routing Support, page 15-1](#)
- [Network Address Translation, page 15-2](#)
- [How Data Moves Through the Security Appliance in Routed Firewall Mode, page 15-3](#)

IP Routing Support

The security appliance acts as a router between connected networks, and each interface requires an IP address on a different subnet. In single context mode, the routed firewall supports OSPF and RIP. Multiple context mode supports static routes only. We recommend using the advanced routing capabilities of the upstream and downstream routers instead of relying on the security appliance for extensive routing needs.

Network Address Translation

NAT substitutes the local address on a packet with a global address that is routable on the destination network. By default, NAT is not required. If you want to enforce a NAT policy that requires hosts on a higher security interface (inside) to use NAT when communicating with a lower security interface (outside), you can enable NAT control (see the **nat-control** command).



Note


NAT control was the default behavior for software versions earlier than Version 7.0. If you upgrade a security appliance from an earlier version, then the **nat-control** command is automatically added to your configuration to maintain the expected behavior.

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.

Figure 15-1 shows a typical NAT scenario, with a private network on the inside. When the inside user sends a packet to a web server on the Internet, the local source address of the packet is changed to a routable global address. When the web server responds, it sends the response to the global address, and the security appliance receives the packet. The security appliance then translates the global address to the local address before sending it on to the user.

Figure 15-1 NAT Example



92405

How Data Moves Through the Security Appliance in Routed Firewall Mode


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

- [An Inside User Visits a Web Server, page 15-3](#)
- [An Outside User Visits a Web Server on the DMZ, page 15-4](#)
- [An Inside User Visits a Web Server on the DMZ, page 15-6](#)
- [An Outside User Attempts to Access an Inside Host, page 15-7](#)
- [A DMZ User Attempts to Access an Inside Host, page 15-8](#)

An Inside User Visits a Web Server

Figure 15-2 shows an inside user accessing an outside web server.

Figure 15-2 *Inside to Outside*



The following steps describe how data moves through the security appliance (see Figure 15-2):

1. The user on the inside network requests a web page from www.example.com.
2. The security appliance receives the packet and because it is a new session, the security appliance verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

For multiple context mode, the security appliance first classifies the packet according to either a unique interface or a unique destination address associated with a context; the destination address is associated by matching an address translation in a context. In this case, the interface would be unique; the `www.example.com` IP address does not have a current address translation in a context.

3. The security appliance translates the local source address (10.1.2.27) to the global address 209.165.201.10, which is on the outside interface subnet.


The global address could be on any subnet, but routing is simplified when it is on the outside interface subnet.

4. The security appliance then records that a session is established and forwards the packet from the outside interface.
5. When `www.example.com` responds to the request, the packet goes through the security appliance, and because the session is already established, the packet bypasses the many lookups associated with a new connection. The security appliance performs NAT by translating the global destination address to the local user address, 10.1.2.27.
6. The security appliance forwards the packet to the inside user.

An Outside User Visits a Web Server on the DMZ

Figure 15-3 shows an outside user accessing the DMZ web server.

Figure 15-3 Outside to DMZ



The following steps describe how data moves through the security appliance (see [Figure 15-3](#)):

1. A user on the outside network requests a web page from the DMZ web server using the global destination address of 209.165.201.3, which is on the outside interface subnet.
2. The security appliance receives the packet and because it is a new session, the security appliance verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).


For multiple context mode, the security appliance first classifies the packet according to either a unique interface or a unique destination address associated with a context; the destination address is associated by matching an address translation in a context. In this case, the classifier “knows” that the DMZ web server address belongs to a certain context because of the server address translation.

3. The security appliance translates the destination address to the local address 10.1.1.3.
4. The security appliance then adds a session entry to the fast path and forwards the packet from the DMZ interface.
5. When the DMZ web server responds to the request, the packet goes through the security appliance and because the session is already established, the packet bypasses the many lookups associated with a new connection. The security appliance performs NAT by translating the local source address to 209.165.201.3.
6. The security appliance forwards the packet to the outside user.

An Inside User Visits a Web Server on the DMZ

Figure 15-4 shows an inside user accessing the DMZ web server.

Figure 15-4 Inside to DMZ



The following steps describe how data moves through the security appliance (see Figure 15-4):

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


For multiple context mode, the security appliance first classifies the packet according to either a unique interface or a unique destination address associated with a context; the destination address is associated by matching an address translation in a context. In this case, the interface is unique; the web server IP address does not have a current address translation.

3. The security appliance then records that a session is established and forwards the packet out of the DMZ interface.
4. When the DMZ web server responds to the request, the packet goes through the fast path, which lets the packet bypass the many lookups associated with a new connection.
5. The security appliance forwards the packet to the inside user.

An Outside User Attempts to Access an Inside Host

Figure 15-5 shows an outside user attempting to access the inside network.

Figure 15-5 *Outside to Inside*




The following steps describe how data moves through the security appliance (see Figure 15-5):

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

A DMZ User Attempts to Access an Inside Host

Figure 15-6 shows a user in the DMZ attempting to access the inside network.

Figure 15-6 DMZ to Inside



The following steps describe how data moves through the security appliance (see Figure 15-6):

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

Transparent Mode Overview

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

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

- [Transparent Firewall Network, page 15-9](#)
- [Allowing Layer 3 Traffic, page 15-9](#)
- [Passing Traffic Not Allowed in Routed Mode, page 15-9](#)
- [MAC Address Lookups, page 15-10](#)
- [Using the Transparent Firewall in Your Network, page 15-10](#)
- [Transparent Firewall Guidelines, page 15-10](#)

- [Unsupported Features in Transparent Mode, page 15-11](#)
- [How Data Moves Through the Transparent Firewall, page 15-13](#)

Transparent Firewall Network

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

Allowing Layer 3 Traffic

IPv4 traffic is allowed through the transparent firewall automatically from a higher security interface to a lower security interface, without an access list. ARPs are allowed through the transparent firewall in both directions without an access list. ARP traffic can be controlled by ARP inspection. For Layer 3 traffic travelling from a low to a high security interface, an extended access list is required.

Allowed MAC Addresses

The following destination MAC addresses are allowed through the transparent firewall. Any MAC address not on this list is dropped.

- TRUE broadcast destination MAC address equal to FFFF.FFFF.FFFF
- IPv4 multicast MAC addresses from 0100.5E00.0000 to 0100.5EFE.FFFF
- IPv6 multicast MAC addresses from 3333.0000.0000 to 3333.FFFF.FFFF
- BPDU multicast address equal to 0100.0CCC.CCCD
- Appletalk multicast MAC addresses from 0900.0700.0000 to 0900.07FF.FFFF

Passing Traffic Not Allowed in Routed Mode

In routed mode, some types of traffic cannot pass through the security appliance even if you allow it in an access list. The transparent firewall, however, can allow almost any traffic through using either an extended access list (for IP traffic) or an EtherType access list (for non-IP traffic).



Note

The transparent mode security appliance does not pass CDP packets or IPv6 packets, or any packets that do not have a valid EtherType greater than or equal to 0x600. For example, you cannot pass IS-IS packets. An exception is made for BPDUs, which are supported.

For example, you can establish routing protocol adjacencies through a transparent firewall; you can allow OSPF, RIP, EIGRP, or BGP traffic through based on an extended access list. Likewise, protocols like HSRP or VRRP can pass through the security appliance.

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

For features that are not directly supported on the transparent firewall, you can allow traffic to pass through so that upstream and downstream routers can support the functionality. For example, by using an extended access list, you can allow DHCP traffic (instead of the unsupported DHCP relay feature) or multicast traffic such as that created by IP/TV.


MAC Address Lookups

When the security appliance runs in transparent mode, the outgoing interface of a packet is determined by performing a MAC address lookup instead of a route lookup. Route statements can still be configured, but they only apply to security appliance-originated traffic. For example, if your syslog server is located on a remote network, you must use a static route so the security appliance can reach that subnet.

Using the Transparent Firewall in Your Network

Figure 15-7 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 15-7 Transparent Firewall Network



92411

Transparent Firewall Guidelines

Follow these guidelines when planning your transparent firewall network:

- A management IP address is required; for multiple context mode, an IP address is required for each context.

Unlike routed mode, which requires an IP address for each interface, a transparent firewall has an IP address assigned to the entire device. The security appliance uses this IP address as the source address for packets originating on the security appliance, such as system messages or AAA communications.

The management IP address must be on the same subnet as the connected network. You cannot set the subnet to a host subnet (255.255.255.255).

You can configure an IP address for the Management 0/0 management-only interface. This IP address can be on a separate subnet from the main management IP address.



Note If the management IP address is not configured, transient traffic does not pass through the transparent firewall. For multiple context mode, transient traffic does not pass through virtual contexts.

- The transparent security appliance uses an inside interface and an outside interface only. If your platform includes a dedicated management interface, you can also configure the management interface or subinterface for management traffic only.
In single mode, you can only use two data interfaces (and the dedicated management interface, if available) even if your security appliance includes more than two interfaces.
- Each directly connected network must be on the same subnet.
- Do not specify the security appliance management IP address as the default gateway for connected devices; devices need to specify the router on the other side of the security appliance as the default gateway.
- For multiple context mode, each context must use different interfaces; you cannot share an interface across contexts.
- For multiple context mode, each context typically uses a different subnet. You can use overlapping subnets, but your network topology requires router and NAT configuration to make it possible from a routing standpoint.

Unsupported Features in Transparent Mode

Table 15-1 lists the features are not supported in transparent mode.

Table 15-1 *Unsupported Features in Transparent Mode*

Feature	Description
Dynamic DNS	—
DHCP relay	The transparent firewall can act as a DHCP server, but it does not support the DHCP relay commands. DHCP relay is not required because you can allow DHCP traffic to pass through using two extended access lists: one that allows DHCP requests from the inside interface to the outside, and one that allows the replies from the server in the other direction.


Table 15-1 *Unsupported Features in Transparent Mode (continued)*

Feature	Description
Dynamic routing protocols	You can, however, add static routes for traffic originating on the security appliance. You can also allow dynamic routing protocols through the security appliance using an extended access list.
IPv6	You also cannot allow IPv6 using an EtherType access list.
Multicast	You can allow multicast traffic through the security appliance by allowing it in an extended access list.
NAT	NAT is performed on the upstream router.
QoS	—
VPN termination for through traffic	The transparent firewall supports site-to-site VPN tunnels for management connections only. It does not terminate VPN connections for traffic through the security appliance. You can pass VPN traffic through the security appliance using an extended access list, but it does not terminate non-management connections. WebVPN is also not supported.

How Data Moves Through the Transparent Firewall

Figure 15-8 shows a typical transparent firewall implementation with an inside network that contains a public web server. The security appliance has an access list so that the inside users can access Internet resources. Another access list lets the outside users access only the web server on the inside network.

Figure 15-8 Typical Transparent Firewall Data Path




This section describes how data moves through the security appliance, and includes the following topics:

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

An Inside User Visits a Web Server

Figure 15-9 shows an inside user accessing an outside web server.

Figure 15-9 *Inside to Outside*



The following steps describe how data moves through the security appliance (see Figure 15-9):

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


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

3. The security appliance records that a session is established.
4. If the destination MAC address is in its table, the security appliance forwards the packet out of the outside interface. The destination MAC address is that of the upstream router, 209.186.201.2.
If the destination MAC address is not in the security appliance table, the security appliance attempts to discover the MAC address by sending an ARP request and a ping. The first packet is dropped.
5. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.
6. The security appliance forwards the packet to the inside user.

An Outside User Visits a Web Server on the Inside Network

Figure 15-10 shows an outside user accessing the inside web server.

Figure 15-10 *Outside to Inside*



The following steps describe how data moves through the security appliance (see Figure 15-10):

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

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


3. The security appliance records that a session is established.
4. If the destination MAC address is in its table, the security appliance forwards the packet out of the inside interface. The destination MAC address is that of the downstream router, 209.186.201.1.
If the destination MAC address is not in the security appliance table, the security appliance attempts to discover the MAC address by sending an ARP request and a ping. The first packet is dropped.
5. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.

6. The security appliance forwards the packet to the outside user.

An Outside User Attempts to Access an Inside Host

Figure 15-11 shows an outside user attempting to access a host on the inside network.

Figure 15-11 *Outside to Inside*



The following steps describe how data moves through the security appliance (see Figure 15-11):

1. A user on the outside network attempts to reach an inside host.
2. The security appliance receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies if the packet is allowed according to the terms of the security policy (access lists, filters, AAA).
For multiple context mode, the security appliance first classifies the packet according to a unique interface.
3. The packet is denied, and the security appliance drops the packet.
4. If the outside user is attempting to attack the inside network, the security appliance employs many technologies to determine if a packet is valid for an already established session.



CHAPTER 16

Identifying Traffic with Access Lists

This chapter describes how to identify traffic with access lists. This chapter includes the following topics:

- [Access List Overview, page 16-1](#)
- [Adding an Extended Access List, page 16-5](#)
- [Adding an EtherType Access List, page 16-8](#)
- [Adding a Standard Access List, page 16-11](#)
- [Adding a Webtype Access List, page 16-11](#)
- [Simplifying Access Lists with Object Grouping, page 16-11](#)
- [Adding Remarks to Access Lists, page 16-18](#)
- [Scheduling Extended Access List Activation, page 16-18](#)
- [Logging Access List Activity, page 16-20](#)

For information about IPv6 access lists, see the [“Configuring IPv6 Access Lists” section on page 12-6](#).

Access List Overview

Access lists are made up of one or more Access Control Entries. An ACE is a single entry in an access list that specifies a permit or deny rule, and is applied to a protocol, a source and destination IP address or network, and optionally the source and destination ports.

Access lists are used in a variety of features. If your feature uses Modular Policy Framework, you can use an access list to identify traffic within a traffic class map. For more information on Modular Policy Framework, see [Chapter 21, “Using Modular Policy Framework.”](#)

This section includes the following topics:

- [Access List Types, page 16-2](#)
- [Access Control Entry Order, page 16-2](#)
- [Access Control Implicit Deny, page 16-3](#)
- [IP Addresses Used for Access Lists When You Use NAT, page 16-3](#)

Access List Types

Table 16-1 lists the types of access lists and some common uses for them.

Table 16-1 Access List Types and Common Uses

Access List Use	Access List Type	Description
Control network access for IP traffic (routed and transparent mode)	Extended	The security appliance does not allow any traffic from a lower security interface to a higher security interface unless it is explicitly permitted by an extended access list. Note To access the security appliance interface for management access, you do not also need an access list allowing the host IP address. You only need to configure management access according to Chapter 40, “Managing System Access.”
Identify traffic for AAA rules	Extended	AAA rules use access lists to identify traffic.
Control network access for IP traffic for a given user	Extended, downloaded from a AAA server per user	You can configure the RADIUS server to download a dynamic access list to be applied to the user, or the server can send the name of an access list that you already configured on the security appliance.
Identify addresses for NAT (policy NAT and NAT exemption)	Extended	Policy NAT lets you identify local traffic for address translation by specifying the source and destination addresses in an extended access list.
Establish VPN access	Extended	You can use an extended access list in VPN commands.
Identify traffic in a traffic class map for Modular Policy Framework	Extended EtherType	Access lists can be used to identify traffic in a class map, which is used for features that support Modular Policy Framework. Features that support Modular Policy Framework include TCP and general connection settings, and inspection.
For transparent firewall mode, control network access for non-IP traffic	EtherType	You can configure an access list that controls traffic based on its EtherType.
Identify OSPF route redistribution	Standard	Standard access lists include only the destination address. You can use a standard access list to control the redistribution of OSPF routes.
Filtering for WebVPN	Webtype	You can configure a Webtype access list to filter URLs.

Access Control Entry Order

An access list is made up of one or more Access Control Entries. Depending on the access list type, you can specify the source and destination addresses, the protocol, the ports (for TCP or UDP), the ICMP type (for ICMP), or the EtherType.

Each ACE that you enter for a given access list name is appended to the end of the access list.

The order of ACEs is important. When the security appliance decides whether to forward or drop a packet, the security appliance tests the packet against each ACE in the order in which the entries are listed. After a match is found, no more ACEs are checked. For example, if you create an ACE at the beginning of an access list that explicitly permits all traffic, no further statements are ever checked.

You can disable an ACE by specifying the keyword **inactive** in the **access-list** command.

Access Control Implicit Deny

Access lists have an implicit deny at the end of the list, so unless you explicitly permit it, traffic cannot pass. For example, if you want to allow all users to access a network through the security appliance except for particular addresses, then you need to deny the particular addresses and then permit all others.


For EtherType access lists, the implicit deny at the end of the access list does not affect IP traffic or ARPs; for example, if you allow EtherType 8037, the implicit deny at the end of the access list does not now block any IP traffic that you previously allowed with an extended access list (or implicitly allowed from a high security interface to a low security interface). However, if you *explicitly* deny all traffic with an EtherType ACE, then IP and ARP traffic is denied.

IP Addresses Used for Access Lists When You Use NAT

When you use NAT, the IP addresses you specify for an access list depend on the interface to which the access list is attached; you need to use addresses that are valid on the network connected to the interface. This guideline applies for both inbound and outbound access lists: the direction does not determine the address used, only the interface does.

For example, you want to apply an access list to the inbound direction of the inside interface. You configure the security appliance to perform NAT on the inside source addresses when they access outside addresses. Because the access list is applied to the inside interface, the source addresses are the original untranslated addresses. Because the outside addresses are not translated, the destination address used in the access list is the real address (see [Figure 16-1](#)).

Figure 16-1 IP Addresses in Access Lists: NAT Used for Source Addresses




See the following commands for this example:

```
hostname(config)# access-list INSIDE extended permit ip 10.1.1.0 255.255.255.0 host
209.165.200.225
```

```
hostname(config)# access-group INSIDE in interface inside
```

If you want to allow an outside host to access an inside host, you can apply an inbound access list on the outside interface. You need to specify the translated address of the inside host in the access list because that address is the address that can be used on the outside network (see [Figure 16-2](#)).

Figure 16-2 IP Addresses in Access Lists: NAT used for Destination Addresses




See the following commands for this example:

```
hostname(config)# access-list OUTSIDE extended permit ip host 209.165.200.225 host
209.165.201.5
hostname(config)# access-group OUTSIDE in interface outside
```


If you perform NAT on both interfaces, keep in mind the addresses that are visible to a given interface. In [Figure 16-3](#), an outside server uses static NAT so that a translated address appears on the inside network.

Figure 16-3 IP Addresses in Access Lists: NAT used for Source and Destination Addresses



See the following commands for this example:

```
hostname(config)# access-list INSIDE extended permit ip 10.1.1.0 255.255.255.0 host
10.1.1.56
hostname(config)# access-group INSIDE in interface inside
```

Adding an Extended Access List

This section describes how to add an extended access list, and includes the following sections:

- [Extended Access List Overview, page 16-5](#)
- [Allowing Broadcast and Multicast Traffic through the Transparent Firewall, page 16-6](#)
- [Adding an Extended ACE, page 16-6](#)

Extended Access List Overview

An extended access list is made up of one or more ACEs, in which you can specify the line number to insert the ACE, source and destination addresses, and, depending on the ACE type, the protocol, the ports (for TCP or UDP), or the ICMP type (for ICMP). You can identify all of these parameters within the **access-list** command, or you can use object groups for each parameter. This section describes how to identify the parameters within the command. To use object groups, see the [“Simplifying Access Lists with Object Grouping”](#) section on page 16-11.

For information about logging options that you can add to the end of the ACE, see the [“Logging Access List Activity” section on page 16-20](#). For information about time range options, see [“Scheduling Extended Access List Activation” section on page 16-18](#).

For TCP and UDP connections, you do not need an access list to allow returning traffic, because the FWSM allows all returning traffic for established, bidirectional connections. For connectionless protocols such as ICMP, however, the security appliance establishes unidirectional sessions, so you either need access lists to allow ICMP in both directions (by applying access lists to the source and destination interfaces), or you need to enable the ICMP inspection engine. The ICMP inspection engine treats ICMP sessions as bidirectional connections.

You can apply only one access list of each type (extended and EtherType) to each direction of an interface. You can apply the same access lists on multiple interfaces. See [Chapter 18, “Permitting or Denying Network Access,”](#) for more information about applying an access list to an interface.

**Note**

If you change the access list configuration, and you do not want to wait for existing connections to time out before the new access list information is used, you can clear the connections using the **clear local-host** command.

Allowing Broadcast and Multicast Traffic through the Transparent Firewall

In routed firewall mode, broadcast and multicast traffic is blocked even if you allow it in an access list, including unsupported dynamic routing protocols and DHCP (unless you configure DHCP relay). Transparent firewall mode can allow any IP traffic through. This feature is especially useful in multiple context mode, which does not allow dynamic routing, for example.

**Note**

Because these special types of traffic are connectionless, you need to apply an extended access list to both interfaces, so returning traffic is allowed through.

[Table 16-2](#) lists common traffic types that you can allow through the transparent firewall.

Table 16-2 *Transparent Firewall Special Traffic*

Traffic Type	Protocol or Port	Notes
DHCP	UDP ports 67 and 68	If you enable the DHCP server, then the security appliance does not pass DHCP packets.
EIGRP	Protocol 88	—
OSPF	Protocol 89	—
Multicast streams	The UDP ports vary depending on the application.	Multicast streams are always destined to a Class D address (224.0.0.0 to 239.x.x.x).
RIP (v1 or v2)	UDP port 520	—

Adding an Extended ACE

When you enter the **access-list** command for a given access list name, the ACE is added to the end of the access list unless you specify the **line** number.

To add an ACE, enter the following command:

```
hostname(config)# access-list access_list_name [line line_number] [extended]
{deny | permit} protocol source_address mask [operator port] dest_address mask
[operator port | icmp_type] [inactive]
```



Tip

Enter the access list name in upper case letters so the name is easy to see in the configuration. You might want to name the access list for the interface (for example, INSIDE), or for the purpose for which it is created (for example, NO_NAT or VPN).

Typically, you identify the **ip** keyword for the protocol, but other protocols are accepted. For a list of protocol names, see the [“Protocols and Applications” section on page D-11](#).

Enter the **host** keyword before the IP address to specify a single address. In this case, do not enter a mask. Enter the **any** keyword instead of the address and mask to specify any address.

You can specify the source and destination ports only for the **tcp** or **udp** protocols. For a list of permitted keywords and well-known port assignments, see the [“TCP and UDP Ports” section on page D-11](#). DNS, Discard, Echo, Ident, NTP, RPC, SUNRPC, and Talk each require one definition for TCP and one for UDP. TACACS+ requires one definition for port 49 on TCP.

Use an *operator* to match port numbers used by the source or destination. The permitted operators are as follows:

- **lt**—less than
- **gt**—greater than
- **eq**—equal to
- **neq**—not equal to
- **range**—an inclusive range of values. When you use this operator, specify two port numbers, for example:

```
range 100 200
```

You can specify the ICMP type only for the **icmp** protocol. Because ICMP is a connectionless protocol, you either need access lists to allow ICMP in both directions (by applying access lists to the source and destination interfaces), or you need to enable the ICMP inspection engine (see the [“Adding an ICMP Type Object Group” section on page 16-15](#)). The ICMP inspection engine treats ICMP sessions as stateful connections. To control ping, specify **echo-reply (0)** (security appliance to host) or **echo (8)** (host to security appliance). See the [“Adding an ICMP Type Object Group” section on page 16-15](#) for a list of ICMP types.

When you specify a network mask, the method is different from the Cisco IOS software **access-list** command. The security appliance uses a network mask (for example, 255.255.255.0 for a Class C mask). The Cisco IOS mask uses wildcard bits (for example, 0.0.0.255).

To make an ACE inactive, use the **inactive** keyword. To reenble it, enter the entire ACE without the **inactive** keyword. This feature lets you keep a record of an inactive ACE in your configuration to make reenabling easier.

To remove an ACE, enter the **no access-list** command with the entire command syntax string as it appears in the configuration:

```
hostname(config)# no access-list access_list_name [line line_number] [extended]
{deny | permit} protocol source_address mask [operator port] dest_address mask
[operator port | icmp_type] [inactive]
```

If the entry that you are removing is the only entry in the access list, the entire access list is removed.

See the following examples:

The following access list allows all hosts (on the interface to which you apply the access list) to go through the security appliance:

```
hostname(config)# access-list ACL_IN extended permit ip any any
```

The following sample access list prevents hosts on 192.168.1.0/24 from accessing the 209.165.201.0/27 network. All other addresses are permitted.

```
hostname(config)# access-list ACL_IN extended deny tcp 192.168.1.0 255.255.255.0
209.165.201.0 255.255.255.224
hostname(config)# access-list ACL_IN extended permit ip any any
```

If you want to restrict access to only some hosts, then enter a limited permit ACE. By default, all other traffic is denied unless explicitly permitted.

```
hostname(config)# access-list ACL_IN extended permit ip 192.168.1.0 255.255.255.0
209.165.201.0 255.255.255.224
```

The following access list restricts all hosts (on the interface to which you apply the access list) from accessing a website at address 209.165.201.29. All other traffic is allowed.

```
hostname(config)# access-list ACL_IN extended deny tcp any host 209.165.201.29 eq www
hostname(config)# access-list ACL_IN extended permit ip any any
```

Adding an EtherType Access List

Transparent firewall mode only

This section describes how to add an EtherType access list, and includes the following sections:

- [EtherType Access List Overview, page 16-8](#)
- [Adding an EtherType ACE, page 16-10](#)

EtherType Access List Overview

An EtherType access list is made up of one or more ACEs that specify an EtherType. This section includes the following topics:

- [Supported EtherTypes, page 16-8](#)
- [Implicit Permit of IP and ARPs Only, page 16-9](#)
- [Implicit and Explicit Deny ACE at the End of an Access List, page 16-9](#)
- [IPv6 Unsupported, page 16-9](#)
- [Using Extended and EtherType Access Lists on the Same Interface, page 16-9](#)
- [Allowing MPLS, page 16-9](#)

Supported EtherTypes

An EtherType ACE controls any EtherType identified by a 16-bit hexadecimal number.

EtherType access lists support Ethernet V2 frames.

802.3-formatted frames are not handled by the access list because they use a length field as opposed to a type field.

BPDUs, which are handled by the access list, are the only exception: they are SNAP-encapsulated, and the security appliance is designed to specifically handle BPDUs.

The security appliance receives trunk port (Cisco proprietary) BPDUs. Trunk BPDUs have VLAN information inside the payload, so the security appliance modifies the payload with the outgoing VLAN if you allow BPDUs.

**Note**

If you use failover, you must allow BPDUs on both interfaces with an EtherType access list to avoid bridging loops.

Implicit Permit of IP and ARPs Only

IPv4 traffic is allowed through the transparent firewall automatically from a higher security interface to a lower security interface, without an access list. ARPs are allowed through the transparent firewall in both directions without an access list. ARP traffic can be controlled by ARP inspection.

However, to allow any traffic with EtherTypes other than IPv4 and ARP, you need to apply an EtherType access list, even from a high security to a low security interface.

Because EtherTypes are connectionless, you need to apply the access list to both interfaces if you want traffic to pass in both directions.

Implicit and Explicit Deny ACE at the End of an Access List

For EtherType access lists, the implicit deny at the end of the access list does not affect IP traffic or ARPs; for example, if you allow EtherType 8037, the implicit deny at the end of the access list does not now block any IP traffic that you previously allowed with an extended access list (or implicitly allowed from a high security interface to a low security interface). However, if you *explicitly* deny all traffic with an EtherType ACE, then IP and ARP traffic is denied.

IPv6 Unsupported

EtherType ACEs do not allow IPv6 traffic, even if you specify the IPv6 EtherType.

Using Extended and EtherType Access Lists on the Same Interface

You can apply only one access list of each type (extended and EtherType) to each direction of an interface. You can also apply the same access lists on multiple interfaces.

Allowing MPLS

If you allow MPLS, ensure that Label Distribution Protocol and Tag Distribution Protocol TCP connections are established through the security appliance by configuring both MPLS routers connected to the security appliance to use the IP address on the security appliance interface as the router-id for LDP or TDP sessions. (LDP and TDP allow MPLS routers to negotiate the labels (addresses) used to forward packets.)

On Cisco IOS routers, enter the appropriate command for your protocol, LDP or TDP. The *interface* is the interface connected to the security appliance.

```
hostname(config)# mpls ldp router-id interface force
```

Or

```
hostname(config)# tag-switching tdp router-id interface force
```

Adding an EtherType ACE

To add an EtherType ACE, enter the following command:

```
hostname(config)# access-list access_list_name ethertype {permit | deny} {ipx | bpdu |
mpls-unicast | mpls-multicast | any | hex_number}
```

The *hex_number* is any EtherType that can be identified by a 16-bit hexadecimal number greater than or equal to 0x600. See RFC 1700, “Assigned Numbers,” at <http://www.ietf.org/rfc/rfc1700.txt> for a list of EtherTypes.

To remove an ACE, enter the **no access-list** command with the entire command syntax string as it appears in the configuration:

```
hostname(config)# no access-list access_list_name [line line_number] [extended]
{deny | permit} protocol source_address mask [operator port] dest_address mask
[operator port | icmp_type] [inactive]
```

To remove an EtherType ACE, enter the **no access-list** command with the entire command syntax string as it appears in the configuration:

```
hostname(config)# no access-list access_list_name ethertype {permit | deny} {ipx | bpdu |
mpls-unicast | mpls-multicast | any | hex_number}
```



Note

If an EtherType access list is configured to **deny all**, all ethernet frames are discarded. Only physical protocol traffic, such as auto-negotiation, is still allowed.

When you enter the **access-list** command for a given access list name, the ACE is added to the end of the access list.



Tip

Enter the *access_list_name* in upper case letters so the name is easy to see in the configuration. You might want to name the access list for the interface (for example, INSIDE), or for the purpose (for example, MPLS or IPX).

For example, the following sample access list allows common EtherTypes originating on the inside interface:

```
hostname(config)# access-list ETHER ethertype permit ipx
hostname(config)# access-list ETHER ethertype permit bpdu
hostname(config)# access-list ETHER ethertype permit mpls-unicast
hostname(config)# access-group ETHER in interface inside
```

The following access list allows some EtherTypes through the security appliance, but denies IPX:

```
hostname(config)# access-list ETHER ethertype deny ipx
hostname(config)# access-list ETHER ethertype permit 0x1234
hostname(config)# access-list ETHER ethertype permit bpdu
hostname(config)# access-list ETHER ethertype permit mpls-unicast
hostname(config)# access-group ETHER in interface inside
hostname(config)# access-group ETHER in interface outside
```

The following access list denies traffic with EtherType 0x1256, but allows all others on both interfaces:

```
hostname(config)# access-list nonIP ethertype deny 1256
hostname(config)# access-list nonIP ethertype permit any
hostname(config)# access-group ETHER in interface inside
hostname(config)# access-group ETHER in interface outside
```

Adding a Standard Access List

Single context mode only

Standard access lists identify the destination IP addresses of OSPF routes, and can be used in a route map for OSPF redistribution. Standard access lists cannot be applied to interfaces to control traffic.

The following command adds a standard ACE. To add another ACE at the end of the access list, enter another **access-list** command specifying the same access list name. Apply the access list using the [“Defining Route Maps” section on page 9-7](#).

To add an ACE, enter the following command:

```
hostname(config)# access-list access_list_name standard {deny | permit} {any | ip_address mask}
```

To remove an ACE, enter the **no access-list** command with the entire command syntax string as it appears in the configuration:

```
hostname(config)# no access-list access_list_name standard {deny | permit} {any | ip_address mask}
```

The following sample access list identifies routes to 192.168.1.0/24:

```
hostname(config)# access-list OSPF standard permit 192.168.1.0 255.255.255.0
```

Adding a Webtype Access List

To add an access list to the configuration that supports filtering for WebVPN, enter the following command:

```
hostname(config)# access-list access_list_name webtype {deny | permit} url [url_string | any]
```

To remove a Webtype access list, enter the **no access-list** command with the entire syntax string as it appears in the configuration:

```
hostname(config)# access-list access_list_name webtype {deny | permit} url [url_string | any]
```

For information about logging options that you can add to the end of the ACE, see the [“Logging Access List Activity” section on page 16-20](#).

Simplifying Access Lists with Object Grouping

This section describes how to use object grouping to simplify access list creation and maintenance.

This section includes the following topics:

- [How Object Grouping Works, page 16-12](#)
- [Adding Object Groups, page 16-12](#)

- [Nesting Object Groups, page 16-15](#)
- [Displaying Object Groups, page 16-17](#)
- [Removing Object Groups, page 16-17](#)
- [Using Object Groups with an Access List, page 16-16](#)

How Object Grouping Works

By grouping like-objects together, you can use the object group in an ACE instead of having to enter an ACE for each object separately. You can create the following types of object groups:

- Protocol
- Network
- Service
- ICMP type

For example, consider the following three object groups:

- **MyServices**—Includes the TCP and UDP port numbers of the service requests that are allowed access to the internal network
- **TrustedHosts**—Includes the host and network addresses allowed access to the greatest range of services and servers
- **PublicServers**—Includes the host addresses of servers to which the greatest access is provided

After creating these groups, you could use a single ACE to allow trusted hosts to make specific service requests to a group of public servers.

You can also nest object groups in other object groups.



Note

The ACE system limit applies to expanded access lists. If you use object groups in ACEs, the number of actual ACEs that you enter is fewer, but the number of expanded ACEs is the same as without object groups. In many cases, object groups create more ACEs than if you added them manually, because creating ACEs manually leads you to summarize addresses more than an object group does. To view the number of expanded ACEs in an access list, enter the **show access-list** *access_list_name* command.

Adding Object Groups

This section describes how to add object groups.

This section includes the following topics:

- [Adding a Protocol Object Group, page 16-13](#)
- [Adding a Network Object Group, page 16-13](#)
- [Adding a Service Object Group, page 16-14](#)
- [Adding an ICMP Type Object Group, page 16-15](#)

Adding a Protocol Object Group

To add or change a protocol object group, follow these steps. After you add the group, you can add more objects as required by following this procedure again for the same group name and specifying additional objects. You do not need to reenter existing objects; the commands you already set remain in place unless you remove them with the **no** form of the command.

To add a protocol group, follow these steps:

-
- Step 1** To add a protocol group, enter the following command:

```
hostname(config)# object-group protocol grp_id
```

The *grp_id* is a text string up to 64 characters in length.

The prompt changes to protocol configuration mode.

- Step 2** (Optional) To add a description, enter the following command:

```
hostname(config-protocol)# description text
```

The description can be up to 200 characters.

- Step 3** To define the protocols in the group, enter the following command for each protocol:

```
hostname(config-protocol)# protocol-object protocol
```

The *protocol* is the numeric identifier of the specific IP protocol (1 to 254) or a keyword identifier (for example, **icmp**, **tcp**, or **udp**). To include all IP protocols, use the keyword **ip**. For a list of protocols you can specify, see the [“Protocols and Applications” section on page D-11](#).

For example, to create a protocol group for TCP, UDP, and ICMP, enter the following commands:

```
hostname(config)# object-group protocol tcp_udp_icmp  
hostname(config-protocol)# protocol-object tcp  
hostname(config-protocol)# protocol-object udp  
hostname(config-protocol)# protocol-object icmp
```

Adding a Network Object Group

To add or change a network object group, follow these steps. After you add the group, you can add more objects as required by following this procedure again for the same group name and specifying additional objects. You do not need to reenter existing objects; the commands you already set remain in place unless you remove them with the **no** form of the command.



Note

A network object group supports IPv4 and IPv6 addresses, depending on the type of access list. For more information about IPv6 access lists, see [“Configuring IPv6 Access Lists” section on page 12-6](#).

To add a network group, follow these steps:

-
- Step 1** To add a network group, enter the following command:

```
hostname(config)# object-group network grp_id
```

The *grp_id* is a text string up to 64 characters in length.

The prompt changes to network configuration mode.

Step 2 (Optional) To add a description, enter the following command:

```
hostname(config-network)# description text
```

The description can be up to 200 characters.

Step 3 To define the networks in the group, enter the following command for each network or address:

```
hostname(config-network)# network-object {host ip_address | ip_address mask}
```

For example, to create network group that includes the IP addresses of three administrators, enter the following commands:

```
hostname(config)# object-group network admins
hostname(config-network)# description Administrator Addresses
hostname(config-network)# network-object host 10.1.1.4
hostname(config-network)# network-object host 10.1.1.78
hostname(config-network)# network-object host 10.1.1.34
```

Adding a Service Object Group

To add or change a service object group, follow these steps. After you add the group, you can add more objects as required by following this procedure again for the same group name and specifying additional objects. You do not need to reenter existing objects; the commands you already set remain in place unless you remove them with the **no** form of the command.

To add a service group, follow these steps:

Step 1 To add a service group, enter the following command:

```
hostname(config)# object-group service grp_id {tcp | udp | tcp-udp}
```

The *grp_id* is a text string up to 64 characters in length.

Specify the protocol for the services (ports) you want to add, either **tcp**, **udp**, or **tcp-udp** keywords. Enter **tcp-udp** keyword if your service uses both TCP and UDP with the same port number, for example, DNS (port 53).

The prompt changes to service configuration mode.

Step 2 (Optional) To add a description, enter the following command:

```
hostname(config-service)# description text
```

The description can be up to 200 characters.

Step 3 To define the ports in the group, enter the following command for each port or range of ports:

```
hostname(config-service)# port-object {eq port | range begin_port end_port}
```

For a list of permitted keywords and well-known port assignments, see the [“Protocols and Applications” section on page D-11](#).

For example, to create service groups that include DNS (TCP/UDP), LDAP (TCP), and RADIUS (UDP), enter the following commands:

```
hostname(config)# object-group service services1 tcp-udp
hostname(config-service)# description DNS Group
hostname(config-service)# port-object eq domain
```

```
hostname(config-service)# object-group service services2 udp
hostname(config-service)# description RADIUS Group
hostname(config-service)# port-object eq radius
hostname(config-service)# port-object eq radius-acct

hostname(config-service)# object-group service services3 tcp
hostname(config-service)# description LDAP Group
hostname(config-service)# port-object eq ldap
```

Adding an ICMP Type Object Group

To add or change an ICMP type object group, follow these steps. After you add the group, you can add more objects as required by following this procedure again for the same group name and specifying additional objects. You do not need to reenter existing objects; the commands you already set remain in place unless you remove them with the **no** form of the command.

To add an ICMP type group, follow these steps:

-
- Step 1** To add an ICMP type group, enter the following command:

```
hostname(config)# object-group icmp-type grp_id
```

The *grp_id* is a text string up to 64 characters in length.

The prompt changes to ICMP type configuration mode.

- Step 2** (Optional) To add a description, enter the following command:

```
hostname(config-icmp-type)# description text
```

The description can be up to 200 characters.

- Step 3** To define the ICMP types in the group, enter the following command for each type:

```
hostname(config-icmp-type)# icmp-object icmp_type
```

See the [“ICMP Types” section on page D-15](#) for a list of ICMP types.

For example, to create an ICMP type group that includes echo-reply and echo (for controlling ping), enter the following commands:

```
hostname(config)# object-group icmp-type ping
hostname(config-service)# description Ping Group
hostname(config-icmp-type)# icmp-object echo
hostname(config-icmp-type)# icmp-object echo-reply
```

Nesting Object Groups

To nest an object group within another object group of the same type, first create the group that you want to nest according to the [“Adding Object Groups” section on page 16-12](#). Then follow these steps:

-
- Step 1** To add or edit an object group under which you want to nest another object group, enter the following command:

```
hostname(config)# object-group {{protocol | network | icmp-type} grp_id | service grp_id
{tcp | udp | tcp-udp}}
```

Step 2 To add the specified group under the object group you specified in Step 1, enter the following command:

```
hostname(config-group_type)# group-object grp_id
```

The nested group must be of the same type.

You can mix and match nested group objects and regular objects within an object group.

For example, you create network object groups for privileged users from various departments:

```
hostname(config)# object-group network eng
hostname(config-network)# network-object host 10.1.1.5
hostname(config-network)# network-object host 10.1.1.9
hostname(config-network)# network-object host 10.1.1.89

hostname(config-network)# object-group network hr
hostname(config-network)# network-object host 10.1.2.8
hostname(config-network)# network-object host 10.1.2.12

hostname(config-network)# object-group network finance
hostname(config-network)# network-object host 10.1.4.89
hostname(config-network)# network-object host 10.1.4.100
```

You then nest all three groups together as follows:

```
hostname(config)# object-group network admin
hostname(config-network)# group-object eng
hostname(config-network)# group-object hr
hostname(config-network)# group-object finance
```

You only need to specify the admin object group in your ACE as follows:

```
hostname(config)# access-list ACL_IN extended permit ip object-group admin host
209.165.201.29
```

Using Object Groups with an Access List

To use object groups in an access list, replace the normal protocol (*protocol*), network (*source_address_mask*, etc.), service (*operator port*), or ICMP type (*icmp_type*) parameter with **object-group grp_id** parameter.

For example, to use object groups for all available parameters in the **access-list {tcp | udp}** command, enter the following command:

```
hostname(config)# access-list access_list_name [line line_number] [extended] {deny /
permit} {tcp | udp} object-group nw_grp_id [object-group svc_grp_id] object-group
nw_grp_id [object-group svc_grp_id] [log [[level] [interval secs] | disable | default]]
[inactive | time-range time_range_name]
```

You do not have to use object groups for all parameters; for example, you can use an object group for the source address, but identify the destination address with an address and mask.

The following normal access list that does not use object groups restricts several hosts on the inside network from accessing several web servers. All other traffic is allowed.

```
hostname(config)# access-list ACL_IN extended deny tcp host 10.1.1.4 host 209.165.201.29
eq www
hostname(config)# access-list ACL_IN extended deny tcp host 10.1.1.78 host 209.165.201.29
eq www
hostname(config)# access-list ACL_IN extended deny tcp host 10.1.1.89 host 209.165.201.29
eq www
```

```

hostname(config)# access-list ACL_IN extended deny tcp host 10.1.1.4 host 209.165.201.16
eq www
hostname(config)# access-list ACL_IN extended deny tcp host 10.1.1.78 host 209.165.201.16
eq www
hostname(config)# access-list ACL_IN extended deny tcp host 10.1.1.89 host 209.165.201.16
eq www
hostname(config)# access-list ACL_IN extended deny tcp host 10.1.1.4 host 209.165.201.78
eq www
hostname(config)# access-list ACL_IN extended deny tcp host 10.1.1.78 host 209.165.201.78
eq www
hostname(config)# access-list ACL_IN extended deny tcp host 10.1.1.89 host 209.165.201.78
eq www
hostname(config)# access-list ACL_IN extended permit ip any any
hostname(config)# access-group ACL_IN in interface inside

```

If you make two network object groups, one for the inside hosts, and one for the web servers, then the configuration can be simplified and can be easily modified to add more hosts:

```

hostname(config)# object-group network denied
hostname(config-network)# network-object host 10.1.1.4
hostname(config-network)# network-object host 10.1.1.78
hostname(config-network)# network-object host 10.1.1.89

hostname(config-network)# object-group network web
hostname(config-network)# network-object host 209.165.201.29
hostname(config-network)# network-object host 209.165.201.16
hostname(config-network)# network-object host 209.165.201.78

hostname(config-network)# access-list ACL_IN extended deny tcp object-group denied
object-group web eq www
hostname(config)# access-list ACL_IN extended permit ip any any
hostname(config)# access-group ACL_IN in interface inside

```

Displaying Object Groups

To display a list of the currently configured object groups, enter the following command:

```
hostname(config)# show object-group [protocol | network | service | icmp-type | id grp_id]
```

If you enter the command without any parameters, the system displays all configured object groups.

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

```

hostname# show object-group
object-group network ftp_servers
  description: This is a group of FTP servers
  network-object host 209.165.201.3
  network-object host 209.165.201.4
object-group network TrustedHosts
  network-object host 209.165.201.1
  network-object 192.168.1.0 255.255.255.0
group-object ftp_servers

```

Removing Object Groups

To remove an object group, enter one of the following commands.



Note

You cannot remove an object group or make an object group empty if it is used in an access list.

- To remove a specific object group, enter the following command:

```
hostname(config)# no object-group grp_id
```

- To remove all object groups of the specified type, enter the following command:

```
hostname(config)# clear object-group [protocol | network | services | icmp-type]
```

If you do not enter a type, all object groups are removed.

Adding Remarks to Access Lists

You can include remarks about entries in any access list, including extended, EtherType, and standard access lists. The remarks make the access list easier to understand.

To add a remark after the last **access-list** command you entered, enter the following command:

```
hostname(config)# access-list access_list_name remark text
```

If you enter the remark before any **access-list** command, then the remark is the first line in the access list.

If you delete an access list using the **no access-list access_list_name** command, then all the remarks are also removed.

The text can be up to 100 characters in length. You can enter leading spaces at the beginning of the text. Trailing spaces are ignored.

For example, you can add remarks before each ACE, and the remark appears in the access list in this location. Entering a dash (-) at the beginning of the remark helps set it apart from ACEs.

```
hostname(config)# access-list OUT remark - this is the inside admin address
hostname(config)# access-list OUT extended permit ip host 209.168.200.3 any
hostname(config)# access-list OUT remark - this is the hr admin address
hostname(config)# access-list OUT extended permit ip host 209.168.200.4 any
```

Scheduling Extended Access List Activation

You can schedule each ACE to be activated at specific times of the day and week by applying a time range to the ACE. This section includes the following topics:

- [Adding a Time Range, page 16-18](#)
- [Applying the Time Range to an ACE, page 16-19](#)

Adding a Time Range

To add a time range to implement a time-based access list, perform the following steps:

-
- Step 1** Identify the time-range name by entering the following command:

```
hostname(config)# time-range name
```

- Step 2** Specify the time range as either a recurring time range or an absolute time range.

**Note**

Users could experience a delay of approximately 80 to 100 seconds after the specified end time for the ACL to become inactive. For example, if the specified end time is 3:50, because the end time is inclusive, the command is picked up anywhere between 3:51:00 and 3:51:59. After the command is picked up, the security appliance finishes any currently running task and then services the command to deactivate the ACL.

Multiple periodic entries are allowed per **time-range** command. If a **time-range** command has both **absolute** and **periodic** values specified, then the **periodic** commands are evaluated only after the **absolute** start time is reached, and are not further evaluated after the **absolute** end time is reached.

- Recurring time range:

```
hostname(config-time-range)# periodic days-of-the-week time to [days-of-the-week] time
```

You can specify the following values for *days-of-the-week*:

- **monday, tuesday, wednesday, thursday, friday, saturday, and sunday.**
- **daily**
- **weekdays**
- **weekend**

The *time* is in the format *hh:mm*. For example, 8:00 is 8:00 a.m. and 20:00 is 8:00 p.m.

- Absolute time range:

```
hostname(config-time-range)# absolute start time date [end time date]
```

The *time* is in the format *hh:mm*. For example, 8:00 is 8:00 a.m. and 20:00 is 8:00 p.m.

The *date* is in the format *day month year*; for example, **1 january 2006**.

The following is an example of an absolute time range beginning at 8:00 a.m. on January 1, 2006. Because no end time and date are specified, the time range is in effect indefinitely.

```
hostname(config)# time-range for2006
hostname(config-time-range)# absolute start 8:00 1 january 2006
```

The following is an example of a weekly periodic time range from 8:00 a.m. to 6:00 p.m. on weekdays.:

```
hostname(config)# time-range workinghours
hostname(config-time-range)# periodic weekdays 8:00 to 18:00
```

Applying the Time Range to an ACE

To apply the time range to an ACE, use the following command:

```
hostname(config)# access-list access_list_name [extended] {deny / permit}...[time-range name]
```

See the [“Adding an Extended Access List” section on page 16-5](#) for complete **access-list** command syntax.

**Note**

If you also enable logging for the ACE, use the **log** keyword before the **time-range** keyword. If you disable the ACE using the **inactive** keyword, use the **inactive** keyword as the last keyword.

The following example binds an access list named “Sales” to a time range named “New_York_Minute.”

```
hostname(config)# access-list Sales line 1 extended deny tcp host 209.165.200.225 host
209.165.201.1 time-range New_York_Minute
```

Logging Access List Activity

This section describes how to configure access list logging for extended access lists and Webtype access lists.

This section includes the following topics:

- [Access List Logging Overview, page 16-20](#)
- [Configuring Logging for an Access Control Entry, page 16-21](#)
- [Managing Deny Flows, page 16-22](#)

Access List Logging Overview

By default, when traffic is denied by an extended ACE or a Webtype ACE, the security appliance generates system message 106023 for each denied packet, in the following form:

```
%ASA|PIX-4-106023: Deny protocol src [interface_name:source_address/source_port] dst
interface_name:dest_address/dest_port [type {string}, code {code}] by access_group acl_id
```

If the security appliance is attacked, the number of system messages for denied packets can be very large. We recommend that you instead enable logging using system message 106100, which provides statistics for each ACE and lets you limit the number of system messages produced. Alternatively, you can disable all logging.

**Note**

Only ACEs in the access list generate logging messages; the implicit deny at the end of the access list does not generate a message. If you want all denied traffic to generate messages, add the implicit ACE manually to the end of the access list, as follows.

```
hostname(config)# access-list TEST deny ip any any log
```

The **log** options at the end of the extended **access-list** command lets you to set the following behavior:

- Enable message 106100 instead of message 106023
- Disable all logging
- Return to the default logging using message 106023

System message 106100 is in the following form:

```
%ASA|PIX-n-106100: access-list acl_id {permitted | denied} protocol
interface_name/source_address(source_port) -> interface_name/dest_address(dest_port)
hit-cnt number ({first hit | number-second interval})
```


When you enable logging for message 106100, if a packet matches an ACE, the security appliance creates a flow entry to track the number of packets received within a specific interval. The security appliance generates a system message at the first hit and at the end of each interval, identifying the total number of hits during the interval. At the end of each interval, the security appliance resets the hit count to 0. If no packets match the ACE during an interval, the security appliance deletes the flow entry.

A flow is defined by the source and destination IP addresses, protocols, and ports. Because the source port might differ for a new connection between the same two hosts, you might not see the same flow increment because a new flow was created for the connection. See the [“Managing Deny Flows” section on page 16-22](#) to limit the number of logging flows.

Permitted packets that belong to established connections do not need to be checked against access lists; only the initial packet is logged and included in the hit count. For connectionless protocols, such as ICMP, all packets are logged even if they are permitted, and all denied packets are logged.

See the *Cisco Security Appliance Logging Configuration and System Log Messages* for detailed information about this system message.

Configuring Logging for an Access Control Entry

To configure logging for an ACE, see the following information about the **log** option:

```
hostname(config)# access-list access_list_name [extended] {deny / permit}...[log [[level]  
[interval secs] | disable | default]]
```

See the [“Adding an Extended Access List” section on page 16-5](#) and [“Adding a Webtype Access List” section on page 16-11](#) for complete **access-list** command syntax.

If you enter the **log** option without any arguments, you enable system log message 106100 at the default level (6) and for the default interval (300 seconds). See the following options:

- **level**—A severity level between 0 and 7. The default is 6.
- **interval secs**—The time interval in seconds between system messages, from 1 to 600. The default is 300. This value is also used as the timeout value for deleting an inactive flow.
- **disable**—Disables all access list logging.
- **default**—Enables logging to message 106023. This setting is the same as having no **log** option.

For example, you configure the following access list:

```
hostname(config)# access-list outside-acl permit ip host 1.1.1.1 any log 7 interval 600  
hostname(config)# access-list outside-acl permit ip host 2.2.2.2 any  
hostname(config)# access-list outside-acl deny ip any any log 2  
hostname(config)# access-group outside-acl in interface outside
```

When a packet is permitted by the first ACE of *outside-acl*, the security appliance generates the following system message:

```
%ASA|PIX-7-106100: access-list outside-acl permitted tcp outside/1.1.1.1(12345) ->  
inside/192.168.1.1(1357) hit-cnt 1 (first hit)
```

Although 20 additional packets for this connection arrive on the outside interface, the traffic does not have to be checked against the access list, and the hit count does not increase.

If one more connection by the same host is initiated within the specified 10 minute interval (and the source and destination ports remain the same), then the hit count is incremented by 1 and the following message is displayed at the end of the 10 minute interval:

```
%ASA|PIX-7-106100: access-list outside-acl permitted tcp outside/1.1.1.1(12345)->  
inside/192.168.1.1(1357) hit-cnt 2 (600-second interval)
```

When a packet is denied by the third ACE, the security appliance generates the following system message:

```
%ASA|PIX-2-106100: access-list outside-acl denied ip outside/3.3.3.3(12345) ->
inside/192.168.1.1(1357) hit-cnt 1 (first hit)
```

20 additional attempts within a 5 minute interval (the default) result in the following message at the end of 5 minutes:

```
%ASA|PIX-2-106100: access-list outside-acl denied ip outside/3.3.3.3(12345) ->
inside/192.168.1.1(1357) hit-cnt 21 (300-second interval)
```

Managing Deny Flows

When you enable logging for message 106100, if a packet matches an ACE, the security appliance creates a flow entry to track the number of packets received within a specific interval. The security appliance has a maximum of 32 K logging flows for ACEs. A large number of flows can exist concurrently at any point of time. To prevent unlimited consumption of memory and CPU resources, the security appliance places a limit on the number of concurrent *deny* flows; the limit is placed only on deny flows (and not permit flows) because they can indicate an attack. When the limit is reached, the security appliance does not create a new deny flow for logging until the existing flows expire.

For example, if someone initiates a DoS attack, the security appliance can create a large number of deny flows in a short period of time. Restricting the number of deny flows prevents unlimited consumption of memory and CPU resources.

When you reach the maximum number of deny flows, the security appliance issues system message 106101:

```
%ASA|PIX-1-106101: The number of ACL log deny-flows has reached limit (number).
```

To configure the maximum number of deny flows and to set the interval between deny flow alert messages (106101), enter the following commands:

- To set the maximum number of deny flows permitted per context before the security appliance stops logging, enter the following command:

```
hostname(config)# access-list deny-flow-max number
```

The *number* is between 1 and 4096. 4096 is the default.

- To set the amount of time between system messages (number 106101) that identify that the maximum number of deny flows was reached, enter the following command:

```
hostname(config)# access-list alert-interval secs
```

The *seconds* are between 1 and 3600. 300 is the default.



CHAPTER 17

Applying NAT

This chapter describes Network Address Translation (NAT). In routed firewall mode, the security appliance can perform NAT between each network.



Note

In transparent firewall mode, the security appliance does not support NAT.

This chapter contains the following sections:

- [NAT Overview, page 17-1](#)
- [Configuring NAT Control, page 17-16](#)
- [Using Dynamic NAT and PAT, page 17-17](#)
- [Using Static NAT, page 17-26](#)
- [Using Static PAT, page 17-27](#)
- [Bypassing NAT, page 17-29](#)
- [NAT Examples, page 17-33](#)

NAT Overview

This section describes how NAT works on the security appliance, and includes the following topics:

- [Introduction to NAT, page 17-2](#)
- [NAT Control, page 17-3](#)
- [NAT Types, page 17-5](#)
- [Policy NAT, page 17-9](#)
- [NAT and Same Security Level Interfaces, page 17-13](#)
- [Order of NAT Commands Used to Match Real Addresses, page 17-14](#)
- [Mapped Address Guidelines, page 17-14](#)
- [DNS and NAT, page 17-14](#)

Introduction to NAT

Address translation substitutes the real address in a packet with a mapped address that is routable on the destination network. NAT is comprised of two steps: the process in which a real address is translated into a mapped address, and then the process to undo translation for returning traffic.

The security appliance translates an address when a NAT rule matches the traffic. If no NAT rule matches, processing for the packet continues. The exception is when you enable NAT control. NAT control requires that packets traversing from a higher security interface (inside) to a lower security interface (outside) match a NAT rule, or else processing for the packet stops. (See the [“Security Level Overview”](#) section on page 7-1 for more information about security levels, and see [“NAT Control”](#) section on page 17-3 for more information about NAT control).

**Note**


In this document, all types of translation are generally referred to as NAT. When discussing NAT, the terms *inside* and *outside* are relative, and represent the security relationship between any two interfaces. The higher security level is inside and the lower security level is outside; for example, interface 1 is at 60 and interface 2 is at 50, so interface 1 is “inside” and interface 2 is “outside.”

Some of the benefits of NAT are as follows:

- You can use private addresses on your inside networks. Private addresses are not routable on the Internet. (See the [“Private Networks”](#) section on page D-2 for more information.)
- NAT hides the real addresses from other networks, so attackers cannot learn the real address of a host.
- You can resolve IP routing problems such as overlapping addresses.

See [Table 25-1 on page 25-3](#) for information about protocols that do not support NAT.

[Figure 17-1](#) shows a typical NAT scenario, with a private network on the inside. When the inside host at 10.1.2.27 sends a packet to a web server, the real source address, 10.1.2.27, of the packet is changed to a mapped address, 209.165.201.10. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the security appliance receives the packet. The security appliance then undoes the translation of the mapped address, 209.165.201.10 back to the real address, 10.1.2.27 before sending it on to the host.


Figure 17-1 NAT Example

See the following commands for this example:

```
hostname(config)# nat (inside) 1 10.1.2.0 255.255.255.0
hostname(config)# global (outside) 1 209.165.201.1-209.165.201.15
```


NAT Control

NAT control requires that packets traversing from an inside interface to an outside interface match a NAT rule; for any host on the inside network to access a host on the outside network, you must configure NAT to translate the inside host address (see [Figure 17-2](#)).

Figure 17-2 NAT Control and Outbound Traffic


Interfaces at the same security level are not required to use NAT to communicate. However, if you configure dynamic NAT or PAT on a same security interface, then all traffic from the interface to a same security interface or an outside interface must match a NAT rule (see [Figure 17-3](#)).

Figure 17-3 NAT Control and Same Security Traffic



Similarly, if you enable outside dynamic NAT or PAT, then all outside traffic must match a NAT rule when it accesses an inside interface (see [Figure 17-4](#)).

Figure 17-4 NAT Control and Inbound Traffic



Static NAT does not cause these restrictions.

By default, NAT control is disabled, so you do not need to perform NAT on any networks unless you choose to perform NAT. If you upgraded from an earlier version of software, however, NAT control might be enabled on your system. Even with NAT control disabled, you need to perform NAT on any addresses for which you configure dynamic NAT. See the [“Dynamic NAT and PAT Implementation” section on page 17-17](#) for more information on how dynamic NAT is applied.

If you want the added security of NAT control but do not want to translate inside addresses in some cases, you can apply a NAT exemption or identity NAT rule on those addresses. (See the [“Bypassing NAT” section on page 17-29](#) for more information).

To configure NAT control, see the [“Configuring NAT Control” section on page 17-16](#).



Note

In multiple context mode, the packet classifier might rely on the NAT configuration to assign packets to contexts if you do not enable unique MAC addresses for shared interfaces. See the [“How the Security Appliance Classifies Packets” section on page 3-3](#) for more information about the relationship between the classifier and NAT.

NAT Types

This section describes the available NAT types. You can implement address translation as dynamic NAT, Port Address Translation, static NAT, or static PAT or as a mix of these types. You can also configure rules to bypass NAT, for example, if you enable NAT control but do not want to perform NAT. This section includes the following topics:

- [Dynamic NAT, page 17-5](#)
- [PAT, page 17-7](#)
- [Static NAT, page 17-7](#)
- [Static PAT, page 17-8](#)
- [Bypassing NAT When NAT Control is Enabled, page 17-9](#)

Dynamic NAT

Dynamic NAT translates a group of real addresses to a pool of mapped addresses that are routable on the destination network. The mapped pool can include fewer addresses than the real group. When a host you want to translate accesses the destination network, the security appliance assigns it an IP address from the mapped pool. The translation is added only when the real host initiates the connection. The translation is in place only for the duration of the connection, and a given user does not keep the same IP address after the translation times out (see the **timeout xlate** command in the *Cisco Security Appliance Command Reference*). Users on the destination network, therefore, cannot reliably initiate a connection to a host that uses dynamic NAT (even if the connection is allowed by an access list), and the security appliance rejects any attempt to connect to a real host address directly. See the following “[Static NAT](#)” or “[Static PAT](#)” sections for reliable access to hosts.

**Note**

In some cases, a translation is added for a connection (see the **show xlate** command) even though the session is denied by the security appliance. This condition occurs with an outbound access list, a management-only interface, or a backup interface. The translation times out normally.

[Figure 17-5](#) shows a remote host attempting to connect to the real address. The connection is denied because the security appliance only allows returning connections to the mapped address.



Figure 17-5 Remote Host Attempts to Connect to the Real Address

Figure 17-6 shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table, so the security appliance drops the packet.

Figure 17-6 Remote Host Attempts to Initiate a Connection to a Mapped Address**Note**

For the duration of the translation, a remote host can initiate a connection to the translated host if an access list allows it. Because the address is unpredictable, a connection to the host is unlikely. However in this case, you can rely on the security of the access list.

Dynamic NAT has these disadvantages:

- If the mapped pool has fewer addresses than the real group, you could run out of addresses if the amount of traffic is more than expected.

Use PAT if this event occurs often, because PAT provides over 64,000 translations using ports of a single address.

- You have to use a large number of routable addresses in the mapped pool; if the destination network requires registered addresses, such as the Internet, you might encounter a shortage of usable addresses.

The advantage of dynamic NAT is that some protocols cannot use PAT. For example, PAT does not work with IP protocols that do not have a port to overload, such as GRE version 0. PAT also does not work with some applications that have a data stream on one port and the control path on another and are not open standard, such as some multimedia applications. See the [“When to Use Application Protocol Inspection” section on page 25-2](#) for more information about NAT and PAT support.

PAT

PAT translates multiple real addresses to a single mapped IP address. Specifically, the security appliance translates the real address and source port (real socket) to the mapped address and a unique port above 1024 (mapped socket). Each connection requires a separate translation, because the source port differs for each connection. For example, 10.1.1.1:1025 requires a separate translation from 10.1.1.1:1026.

After the connection expires, the port translation also expires after 30 seconds of inactivity. The timeout is not configurable. Users on the destination network cannot reliably initiate a connection to a host that uses PAT (even if the connection is allowed by an access list). Not only can you not predict the real or mapped port number of the host, but the security appliance does not create a translation at all unless the translated host is the initiator. See the following [“Static NAT”](#) or [“Static PAT”](#) sections for reliable access to hosts.

PAT lets you use a single mapped address, thus conserving routable addresses. You can even use the security appliance interface IP address as the PAT address. PAT does not work with some multimedia applications that have a data stream that is different from the control path. See the [“When to Use Application Protocol Inspection” section on page 25-2](#) for more information about NAT and PAT support.



Note

For the duration of the translation, a remote host can initiate a connection to the translated host if an access list allows it. Because the port address (both real and mapped) is unpredictable, a connection to the host is unlikely. Nevertheless, in this case, you can rely on the security of the access list. However, policy PAT does not support time-based ACLs.

Static NAT

Static NAT creates a fixed translation of real address(es) to mapped address(es). With dynamic NAT and PAT, each host uses a different address or port for each subsequent translation. Because the mapped address is the same for each consecutive connection with static NAT, and a persistent translation rule exists, static NAT allows hosts on the destination network to initiate traffic to a translated host (if there is an access list that allows it).

The main difference between dynamic NAT and a range of addresses for static NAT is that static NAT allows a remote host to initiate a connection to a translated host (if there is an access list that allows it), while dynamic NAT does not. You also need an equal number of mapped addresses as real addresses with static NAT.

Static PAT


Static PAT is the same as static NAT, except it lets you specify the protocol (TCP or UDP) and port for the real and mapped addresses.

This feature lets you identify the same mapped address across many different static statements, so long as the port is different for each statement (you cannot use the same mapped address for multiple static NAT statements).

For applications that require application inspection for secondary channels (FTP, VoIP, etc.), the security appliance automatically translates the secondary ports.

For example, if you want to provide a single address for remote users to access FTP, HTTP, and SMTP, but these are all actually different servers on the real network, you can specify static PAT statements for each server that uses the same mapped IP address, but different ports (see Figure 17-7).

Figure 17-7 Static PAT



See the following commands for this example:

```
hostname(config)# static (inside,outside) tcp 209.165.201.3 ftp 10.1.2.27 ftp netmask
255.255.255.255
hostname(config)# static (inside,outside) tcp 209.165.201.3 http 10.1.2.28 http netmask
255.255.255.255
hostname(config)# static (inside,outside) tcp 209.165.201.3 smtp 10.1.2.29 smtp netmask
255.255.255.255
```

You can also use static PAT to translate a well-known port to a non-standard port or vice versa. For example, if your inside web servers use port 8080, you can allow outside users to connect to port 80, and then undo translation to the original port 8080. Similarly, if you want to provide extra security, you can tell your web users to connect to non-standard port 6785, and then undo translation to port 80.

Bypassing NAT When NAT Control is Enabled

If you enable NAT control, then inside hosts must match a NAT rule when accessing outside hosts. If you do not want to perform NAT for some hosts, then you can bypass NAT for those hosts (alternatively, you can disable NAT control). You might want to bypass NAT, for example, if you are using an application that does not support NAT (see the [“When to Use Application Protocol Inspection”](#) section on page 25-2 for information about inspection engines that do not support NAT).

You can configure traffic to bypass NAT using one of three methods. All methods achieve compatibility with inspection engines. However, each method offers slightly different capabilities, as follows:

- Identity NAT (**nat 0** command)—When you configure identity NAT (which is similar to dynamic NAT), you do not limit translation for a host on specific interfaces; you must use identity NAT for connections through all interfaces. Therefore, you cannot choose to perform normal translation on real addresses when you access interface A, but use identity NAT when accessing interface B. Regular dynamic NAT, on the other hand, lets you specify a particular interface on which to translate the addresses. Make sure that the real addresses for which you use identity NAT are routable on all networks that are available according to your access lists.

For identity NAT, even though the mapped address is the same as the real address, you cannot initiate a connection from the outside to the inside (even if the interface access list allows it). Use static identity NAT or NAT exemption for this functionality.

- Static identity NAT (**static** command)—Static identity NAT lets you specify the interface on which you want to allow the real addresses to appear, so you can use identity NAT when you access interface A, and use regular translation when you access interface B. Static identity NAT also lets you use policy NAT, which identifies the real and destination addresses when determining the real addresses to translate (see the [“Policy NAT”](#) section on page 17-9 for more information about policy NAT). For example, you can use static identity NAT for an inside address when it accesses the outside interface and the destination is server A, but use a normal translation when accessing the outside server B.
- NAT exemption (**nat 0 access-list** command)—NAT exemption allows both translated and remote hosts to initiate connections. Like identity NAT, you do not limit translation for a host on specific interfaces; you must use NAT exemption for connections through all interfaces. However, NAT exemption does let you specify the real and destination addresses when determining the real addresses to translate (similar to policy NAT), so you have greater control using NAT exemption. However unlike policy NAT, NAT exemption does not consider the ports in the access list.

Policy NAT

Policy NAT lets you identify real addresses for address translation by specifying the source and destination addresses in an extended access list. You can also optionally specify the source and destination ports. Regular NAT can only consider the real addresses. For example, you can use translate the real address to mapped address A when it accesses server A, but translate the real address to mapped address B when it accesses server B.

**Note**

Policy NAT does not support time-based ACLs.


When you specify the ports in policy NAT for applications that require application inspection for secondary channels (FTP, VoIP, etc.), the security appliance automatically translates the secondary ports.

**Note**

All types of NAT support policy NAT except for NAT exemption. NAT exemption uses an access list to identify the real addresses, but differs from policy NAT in that the ports are not considered. See the [“Bypassing NAT” section on page 17-29](#) for other differences. You can accomplish the same result as NAT exemption using static identity NAT, which does support policy NAT.

Figure 17-8 shows a host on the 10.1.2.0/24 network accessing two different servers. When the host accesses the server at 209.165.201.11, the real address is translated to 209.165.202.129. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130 so that the host appears to be on the same network as the servers, which can help with routing.

Figure 17-8 Policy NAT with Different Destination Addresses




See the following commands for this example:

```
hostname(config)# access-list NET1 permit ip 10.1.2.0 255.255.255.0 209.165.201.0
255.255.255.224
hostname(config)# access-list NET2 permit ip 10.1.2.0 255.255.255.0 209.165.200.224
255.255.255.224
hostname(config)# nat (inside) 1 access-list NET1
hostname(config)# global (outside) 1 209.165.202.129
hostname(config)# nat (inside) 2 access-list NET2
hostname(config)# global (outside) 2 209.165.202.130
```

Figure 17-9 shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for web services, the real address is translated to 209.165.202.129. When the host accesses the same server for Telnet services, the real address is translated to 209.165.202.130.

Figure 17-9 Policy NAT with Different Destination Ports




See the following commands for this example:

```
hostname(config)# access-list WEB permit tcp 10.1.2.0 255.255.255.0 209.165.201.11
255.255.255.255 eq 80
hostname(config)# access-list TELNET permit tcp 10.1.2.0 255.255.255.0 209.165.201.11
255.255.255.255 eq 23
hostname(config)# nat (inside) 1 access-list WEB
hostname(config)# global (outside) 1 209.165.202.129
hostname(config)# nat (inside) 2 access-list TELNET
hostname(config)# global (outside) 2 209.165.202.130
```

For policy static NAT (and for NAT exemption, which also uses an access list to identify traffic), both translated and remote hosts can originate traffic. For traffic originated on the translated network, the NAT access list specifies the real addresses and the *destination* addresses, but for traffic originated on the remote network, the access list identifies the real addresses and the *source* addresses of remote hosts who are allowed to connect to the host using this translation.

Figure 17-10 shows a remote host connecting to a translated host. The translated host has a policy static NAT translation that translates the real address only for traffic to and from the 209.165.201.0/27 network. A translation does not exist for the 209.165.200.224/27 network, so the translated host cannot connect to that network, nor can a host on that network connect to the translated host.

Figure 17-10 Policy Static NAT with Destination Address Translation



See the following commands for this example:

```
hostname(config)# access-list NET1 permit ip 10.1.2.0 255.255.255.224 209.165.201.0 255.255.255.224
hostname(config)# static (inside,outside) 209.165.202.128 access-list NET1
```



Note


For policy static NAT, in undoing the translation, the ACL in the **static** command is not used. If the destination address in the packet matches the mapped address in the static rule, the static rule is used to untranslate the address.



Note

Policy NAT does not support SQL*Net, but it is supported by regular NAT. See the [“When to Use Application Protocol Inspection”](#) section on page 25-2 for information about NAT support for other protocols.

You cannot use policy static NAT to translate different real addresses to the same mapped address. For example, Figure 17-11 shows two inside hosts, 10.1.1.1 and 10.1.1.2, that you want to be translated to 209.165.200.225. When outside host 209.165.201.1 connects to 209.165.200.225, then the connection goes to 10.1.1.1. When outside host 209.165.201.2 connects to the same mapped address, 209.165.200.225, you want the connection to go to 10.1.1.2. However, only one source address in the access list can be used. Since the first ACE is for 10.1.1.1, then all inbound connections sourced from 209.165.201.1 and 209.165.201.2 and destined to 209.165.200.225 will have their destination address translated to 10.1.1.1.

Figure 17-11 Real Addresses Cannot Share the Same Mapped Address

See the following commands for this example. (Although the second ACE in the example does allow 209.165.201.2 to connect to 209.165.200.225, it only allows 209.165.200.225 to be translated to 10.1.1.1.)

```
hostname(config)# static (in,out) 209.165.200.225 access-list policy-nat
hostname(config)# access-list policy-nat permit ip host 10.1.1.1 host 209.165.201.1
hostname(config)# access-list policy-nat permit ip host 10.1.1.2 host 209.165.201.2
```

NAT and Same Security Level Interfaces

NAT is not required between same security level interfaces even if you enable NAT control. You can optionally configure NAT if desired. However, if you configure dynamic NAT when NAT control is enabled, then NAT is required. See the [“NAT Control” section on page 17-3](#) for more information. Also, when you specify a group of IP address(es) for dynamic NAT or PAT on a same security interface, then you must perform NAT on that group of addresses when they access any lower or same security level interface (even when NAT control is not enabled). Traffic identified for static NAT is not affected.

See the [“Allowing Communication Between Interfaces on the Same Security Level” section on page 7-6](#) to enable same security communication.



Note

The security appliance does not support VoIP inspection engines when you configure NAT on same security interfaces. These inspection engines include Skinny, SIP, and H.323. See the [“When to Use Application Protocol Inspection” section on page 25-2](#) for supported inspection engines.

Order of NAT Commands Used to Match Real Addresses

The security appliance matches real addresses to NAT commands in the following order:

1. NAT exemption (**nat 0 access-list**)—In order, until the first match. Identity NAT is not included in this category; it is included in the regular static NAT or regular NAT category. We do not recommend overlapping addresses in NAT exemption statements because unexpected results can occur.
2. Static NAT and Static PAT (regular and policy) (**static**)—In order, until the first match. Static identity NAT is included in this category.
3. Policy dynamic NAT (**nat access-list**)—In order, until the first match. Overlapping addresses are allowed.
4. Regular dynamic NAT (**nat**)—Best match. Regular identity NAT is included in this category. The order of the NAT commands does not matter; the NAT statement that best matches the real address is used. For example, you can create a general statement to translate all addresses (0.0.0.0) on an interface. If you want to translate a subset of your network (10.1.1.1) to a different address, then you can create a statement to translate only 10.1.1.1. When 10.1.1.1 makes a connection, the specific statement for 10.1.1.1 is used because it matches the real address best. We do not recommend using overlapping statements; they use more memory and can slow the performance of the security appliance.

Mapped Address Guidelines

When you translate the real address to a mapped address, you can use the following mapped addresses:

- Addresses on the same network as the mapped interface.

If you use addresses on the same network as the mapped interface (through which traffic exits the security appliance), the security appliance uses proxy ARP to answer any requests for mapped addresses, and thus intercepts traffic destined for a real address. This solution simplifies routing, because the security appliance does not have to be the gateway for any additional networks. However, this approach does put a limit on the number of available addresses used for translations.

For PAT, you can even use the IP address of the mapped interface.

- Addresses on a unique network.

If you need more addresses than are available on the mapped interface network, you can identify addresses on a different subnet. The security appliance uses proxy ARP to answer any requests for mapped addresses, and thus intercepts traffic destined for a real address. If you use OSPF, and you advertise routes on the mapped interface, then the security appliance advertises the mapped addresses. If the mapped interface is passive (not advertising routes) or you are using static routing, then you need to add a static route on the upstream router that sends traffic destined for the mapped addresses to the security appliance.

DNS and NAT


You might need to configure the security appliance to modify DNS replies by replacing the address in the reply with an address that matches the NAT configuration. You can configure DNS modification when you configure each translation.

For example, a DNS server is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure the security appliance to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network (see

Figure 17-12). In this case, you want to enable DNS reply modification on this static statement so that inside users who have access to ftp.cisco.com using the real address receive the real address from the DNS server, and not the mapped address.

When an inside host sends a DNS request for the address of ftp.cisco.com, the DNS server replies with the mapped address (209.165.201.10). The security appliance refers to the static statement for the inside server and translates the address inside the DNS reply to 10.1.3.14. If you do not enable DNS reply modification, then the inside host attempts to send traffic to 209.165.201.10 instead of accessing ftp.cisco.com directly.

Figure 17-12 DNS Reply Modification



See the following command for this example:

```
hostname(config)# static (inside,outside) 209.165.201.10 10.1.3.14 netmask 255.255.255.255 dns
```




Note

If a user on a different network (for example, DMZ) also requests the IP address for ftp.cisco.com from the outside DNS server, then the IP address in the DNS reply is also modified for this user, even though the user is not on the Inside interface referenced by the **static** command.

Figure 17-13 shows a web server and DNS server on the outside. The security appliance has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.201.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.

Figure 17-13 DNS Reply Modification Using Outside NAT



See the following command for this example:

```
hostname(config)# static (outside,inside) 10.1.2.56 209.165.201.10 netmask 255.255.255.255 dns
```

Configuring NAT Control

NAT control requires that packets traversing from an inside interface to an outside interface match a NAT rule. See the “NAT Control” section on page 17-3 for more information.

To enable NAT control, enter the following command:

```
hostname(config)# nat-control
```

To disable NAT control, enter the **no** form of the command.

Using Dynamic NAT and PAT


This section describes how to configure dynamic NAT and PAT, and includes the following topics:

- [Dynamic NAT and PAT Implementation, page 17-17](#)
- [Configuring Dynamic NAT or PAT, page 17-23](#)

Dynamic NAT and PAT Implementation

For dynamic NAT and PAT, you first configure a **nat** command identifying the real addresses on a given interface that you want to translate. Then you configure a separate **global** command to specify the mapped addresses when exiting another interface (in the case of PAT, this is one address). Each **nat** command matches a **global** command by comparing the NAT ID, a number that you assign to each command (see [Figure 17-14](#)).

Figure 17-14 *nat and global ID Matching*




See the following commands for this example:

```
hostname(config)# nat (inside) 1 10.1.2.0 255.255.255.0
hostname(config)# global (outside) 1 209.165.201.3-209.165.201.10
```

You can enter a **nat** command for each interface using the same NAT ID; they all use the same **global** command when traffic exits a given interface. For example, you can configure **nat** commands for Inside and DMZ interfaces, both on NAT ID 1. Then you configure a **global** command on the Outside interface that is also on ID 1. Traffic from the Inside interface and the DMZ interface share a mapped pool or a PAT address when exiting the Outside interface (see Figure 17-15).

Figure 17-15 *nat Commands on Multiple Interfaces*




See the following commands for this example:

```
hostname(config)# nat (inside) 1 10.1.2.0 255.255.255.0
hostname(config)# nat (dmz) 1 10.1.1.0 255.255.255.0
hostname(config)# global (outside) 1 209.165.201.3-209.165.201.10
```

You can also enter a **global** command for each interface using the same NAT ID. If you enter a **global** command for the Outside and DMZ interfaces on ID 1, then the Inside **nat** command identifies traffic to be translated when going to both the Outside and the DMZ interfaces. Similarly, if you also enter a **nat** command for the DMZ interface on ID 1, then the **global** command on the Outside interface is also used for DMZ traffic. (See Figure 17-16).

Figure 17-16 *global and nat Commands on Multiple Interfaces*




See the following commands for this example:

```
hostname(config)# nat (inside) 1 10.1.2.0 255.255.255.0
hostname(config)# nat (dmz) 1 10.1.1.0 255.255.255.0
hostname(config)# global (outside) 1 209.165.201.3-209.165.201.10
hostname(config)# global (dmz) 1 10.1.1.23
```

If you use different NAT IDs, you can identify different sets of real addresses to have different mapped addresses. For example, on the Inside interface, you can have two **nat** commands on two different NAT IDs. On the Outside interface, you configure two **global** commands for these two IDs. Then, when traffic from Inside network A exits the Outside interface, the IP addresses are translated to pool A addresses; while traffic from Inside network B are translated to pool B addresses (see Figure 17-17). If you use policy NAT, you can specify the same real addresses for multiple **nat** commands, as long as the the destination addresses and ports are unique in each access list.

Figure 17-17 Different NAT IDs




See the following commands for this example:

```
hostname(config)# nat (inside) 1 10.1.2.0 255.255.255.0
hostname(config)# nat (inside) 2 192.168.1.0 255.255.255.0
hostname(config)# global (outside) 1 209.165.201.3-209.165.201.10
hostname(config)# global (outside) 2 209.165.201.11
```

You can enter multiple **global** commands for one interface using the same NAT ID; the security appliance uses the dynamic NAT **global** commands first, in the order they are in the configuration, and then uses the PAT **global** commands in order. You might want to enter both a dynamic NAT **global** command and a PAT **global** command if you need to use dynamic NAT for a particular application, but want to have a backup PAT statement in case all the dynamic NAT addresses are depleted. Similarly, you might enter two PAT statements if you need more than the approximately 64,000 PAT sessions that a single PAT mapped statement supports (see [Figure 17-18](#)).

Figure 17-18 NAT and PAT Together




See the following commands for this example:

```
hostname(config)# nat (inside) 1 10.1.2.0 255.255.255.0
hostname(config)# global (outside) 1 209.165.201.3-209.165.201.4
hostname(config)# global (outside) 1 209.165.201.5
```

For outside NAT, you need to identify the **nat** command for outside NAT (the **outside** keyword). If you also want to translate the same traffic when it accesses an inside interface (for example, traffic on a DMZ is translated when accessing the Inside and the Outside interfaces), then you must configure a separate **nat** command without the **outside** option. In this case, you can identify the same addresses in both statements and use the same NAT ID (see Figure 17-19). Note that for outside NAT (DMZ interface to Inside interface), the inside host uses a **static** command to allow outside access, so both the source and destination addresses are translated.

Figure 17-19 Outside NAT and Inside NAT Combined



See the following commands for this example:

```
hostname(config)# nat (dmz) 1 10.1.1.0 255.255.255.0 outside
hostname(config)# nat (dmz) 1 10.1.1.0 255.255.255.0
hostname(config)# static (inside,dmz) 10.1.1.5 10.1.2.27 netmask 255.255.255.255
hostname(config)# global (outside) 1 209.165.201.3-209.165.201.4
hostname(config)# global (inside) 1 10.1.2.30-1-10.1.2.40
```

When you specify a group of IP address(es) in a **nat** command, then you must perform NAT on that group of addresses when they access any lower or same security level interface; you must apply a **global** command with the same NAT ID on each interface, or use a **static** command. NAT is not required for that group when it accesses a higher security interface, because to perform NAT from outside to inside, you must create a separate **nat** command using the **outside** keyword. If you do apply outside NAT, then the NAT requirements preceding come into effect for that group of addresses when they access all higher security interfaces. Traffic identified by a **static** command is not affected.

Configuring Dynamic NAT or PAT

This section describes how to configure dynamic NAT or dynamic PAT. The configuration for dynamic NAT and PAT are almost identical; for NAT you specify a range of mapped addresses, and for PAT you specify a single address.

Figure 17-20 shows a typical dynamic NAT scenario. Only translated hosts can create a NAT session, and responding traffic is allowed back. The mapped address is dynamically assigned from a pool defined by the **global** command.

Figure 17-20 *Dynamic NAT*





Figure 17-21 shows a typical dynamic PAT scenario. Only translated hosts can create a NAT session, and responding traffic is allowed back. The mapped address defined by the **global** command is the same for each translation, but the port is dynamically assigned.

Figure 17-21 *Dynamic PAT*



For more information about dynamic NAT, see the “[Dynamic NAT](#)” section on page 17-5. For more information about PAT, see the “[PAT](#)” section on page 17-7.



Note

If you change the NAT configuration, and you do not want to wait for existing translations to time out before the new NAT information is used, you can clear the translation table using the **clear xlate** command. However, clearing the translation table disconnects all current connections that use translations.

To configure dynamic NAT or PAT, perform the following steps:

Step 1 To identify the real addresses that you want to translate, enter one of the following commands:

- Policy NAT:

```
hostname(config)# nat (real_interface) nat_id access-list acl_name [dns] [outside]
[norandomseq] [[tcp] tcp_max_conns [emb_limit]] [udp udp_max_conns]
```

You can identify overlapping addresses in other **nat** commands. For example, you can identify 10.1.1.0 in one command, but 10.1.1.1 in another. The traffic is matched to a policy NAT command in order, until the first match, or for regular NAT, using the best match.

See the following description about options for this command:

- **access-list** *acl_name*—Identify the real addresses and destination addresses using an extended access list. Create the access list using the **access-list** command (see the [“Adding an Extended Access List”](#) section on page 16-5). This access list should include only **permit** ACEs. You can optionally specify the real and destination ports in the access list using the **eq** operator. Policy NAT considers the **inactive** and **time-range** keywords, but it does not support ACL with all **inactive** and **time-range** ACEs.
- **nat_id**—An integer between 1 and 65535. The NAT ID should match a **global** command NAT ID. See the [“Dynamic NAT and PAT Implementation”](#) section on page 17-17 for more information about how NAT IDs are used. **0** is reserved for NAT exemption. (See the [“Configuring NAT Exemption”](#) section on page 17-32 for more information about NAT exemption.)
- **dns**—If your **nat** command includes the address of a host that has an entry in a DNS server, and the DNS server is on a different interface from a client, then the client and the DNS server need different addresses for the host; one needs the mapped address and one needs the real address. This option rewrites the address in the DNS reply to the client. The translated host needs to be on the same interface as either the client or the DNS server. Typically, hosts that need to allow access from other interfaces use a static translation, so this option is more likely to be used with the **static** command. (See the [“DNS and NAT”](#) section on page 17-14 for more information.)
- **outside**—If this interface is on a lower security level than the interface you identify by the matching **global** statement, then you must enter **outside** to identify the NAT instance as outside NAT.
- **norandomseq**, **tcp** *tcp_max_conns*, **udp** *udp_max_conns*, and *emb_limit*—These keywords set connection limits. However, we recommend using a more versatile method for setting connection limits; see the [“Configuring Connection Limits and Timeouts”](#) section on page 23-6.
- Regular NAT:

```
hostname(config)# nat (real_interface) nat_id real_ip [mask] [dns] [outside]
[norandomseq] [[tcp] tcp_max_conns [emb_limit]] [udp udp_max_conns]
```

The *nat_id* is an integer between 1 and 2147483647. The NAT ID must match a **global** command NAT ID. See the [“Dynamic NAT and PAT Implementation”](#) section on page 17-17 for more information about how NAT IDs are used. **0** is reserved for identity NAT. See the [“Configuring Identity NAT”](#) section on page 17-30 for more information about identity NAT.

See the preceding policy NAT command for information about other options.

Step 2 To identify the mapped address(es) to which you want to translate the real addresses when they exit a particular interface, enter the following command:

```
hostname(config)# global (mapped_interface) nat_id {mapped_ip[-mapped_ip] | interface}
```

This NAT ID should match a **nat** command NAT ID. The matching **nat** command identifies the addresses that you want to translate when they exit this interface.

You can specify a single address (for PAT) or a range of addresses (for NAT). The range can go across subnet boundaries if desired. For example, you can specify the following “supernet”:

192.168.1.1-192.168.2.254

For example, to translate the 10.1.1.0/24 network on the inside interface, enter the following command:

```
hostname(config)# nat (inside) 1 10.1.1.0 255.255.255.0
hostname(config)# global (outside) 1 209.165.201.1-209.165.201.30
```

To identify a pool of addresses for dynamic NAT as well as a PAT address for when the NAT pool is exhausted, enter the following commands:

```
hostname(config)# nat (inside) 1 10.1.1.0 255.255.255.0
hostname(config)# global (outside) 1 209.165.201.5
hostname(config)# global (outside) 1 209.165.201.10-209.165.201.20
```

To translate the lower security dmz network addresses so they appear to be on the same network as the inside network (10.1.1.0), for example, to simplify routing, enter the following commands:

```
hostname(config)# nat (dmz) 1 10.1.2.0 255.255.255.0 outside dns
hostname(config)# global (inside) 1 10.1.1.45
```

To identify a single real address with two different destination addresses using policy NAT, enter the following commands (see [Figure 17-8 on page 17-10](#) for a related figure):

```
hostname(config)# access-list NET1 permit ip 10.1.2.0 255.255.255.0 209.165.201.0
255.255.255.224
hostname(config)# access-list NET2 permit ip 10.1.2.0 255.255.255.0 209.165.200.224
255.255.255.224
hostname(config)# nat (inside) 1 access-list NET1 tcp 0 2000 udp 10000
hostname(config)# global (outside) 1 209.165.202.129
hostname(config)# nat (inside) 2 access-list NET2 tcp 1000 500 udp 2000
hostname(config)# global (outside) 2 209.165.202.130
```

To identify a single real address/destination address pair that use different ports using policy NAT, enter the following commands (see [Figure 17-9 on page 17-11](#) for a related figure):


```
hostname(config)# access-list WEB permit tcp 10.1.2.0 255.255.255.0 209.165.201.11
255.255.255.255 eq 80
hostname(config)# access-list TELNET permit tcp 10.1.2.0 255.255.255.0 209.165.201.11
255.255.255.255 eq 23
hostname(config)# nat (inside) 1 access-list WEB
hostname(config)# global (outside) 1 209.165.202.129
hostname(config)# nat (inside) 2 access-list TELNET
hostname(config)# global (outside) 2 209.165.202.130
```

Using Static NAT

This section describes how to configure a static translation.

Figure 17-22 shows a typical static NAT scenario. The translation is always active so both translated and remote hosts can originate connections, and the mapped address is statically assigned by the **static** command.

Figure 17-22 Static NAT



You cannot use the same real or mapped address in multiple **static** commands between the same two interfaces. Do not use a mapped address in the **static** command that is also defined in a **global** command for the same mapped interface.

For more information about static NAT, see the “Static NAT” section on page 17-7.



Note

If you remove a **static** command, existing connections that use the translation are not affected. To remove these connections, enter the **clear local-host** command.

You cannot clear static translations from the translation table with the **clear xlate** command; you must remove the **static** command instead. Only dynamic translations created by the **nat** and **global** commands can be removed with the **clear xlate** command.

To configure static NAT, enter one of the following commands.

- For policy static NAT, enter the following command:

```
hostname(config)# static (real_interface,mapped_interface) {mapped_ip | interface}
access-list acl_name [dns] [norandomseq] [[tcp] tcp_max_conns [emb_limit]]
[udp udp_max_conns]
```

Create the access list using the **access-list** command (see the “Adding an Extended Access List” section on page 16-5). This access list should include only **permit** ACEs. The source subnet mask used in the access list is also used for the mapped addresses. You can also specify the real and destination ports in the access list using the **eq** operator. Policy NAT does not consider the **inactive** or **time-range** keywords; all ACEs are considered to be active for policy NAT configuration. See the “Policy NAT” section on page 17-9 for more information.

If you specify a network for translation (for example, 10.1.1.0 255.255.255.0), then the security appliance translates the .0 and .255 addresses. If you want to prevent access to these addresses, be sure to configure an access list to deny access.

See the “Configuring Dynamic NAT or PAT” section on page 17-23 for information about the other options.

- To configure regular static NAT, enter the following command:

```
hostname(config)# static (real_interface,mapped_interface) {mapped_ip | interface}
real_ip [netmask mask] [dns] [norandomseq] [[tcp] tcp_max_conns [emb_limit]]
[udp udp_max_conns]
```

See the “Configuring Dynamic NAT or PAT” section on page 17-23 for information about the options.

For example, the following policy static NAT example shows a single real address that is translated to two mapped addresses depending on the destination address (see Figure 17-8 on page 17-10 for a related figure):

```
hostname(config)# access-list NET1 permit ip host 10.1.2.27 209.165.201.0 255.255.255.224
hostname(config)# access-list NET2 permit ip host 10.1.2.27 209.165.200.224
255.255.255.224
hostname(config)# static (inside,outside) 209.165.202.129 access-list NET1
hostname(config)# static (inside,outside) 209.165.202.130 access-list NET2
```

The following command maps an inside IP address (10.1.1.3) to an outside IP address (209.165.201.12):

```
hostname(config)# static (inside,outside) 209.165.201.12 10.1.1.3 netmask 255.255.255.255
```

The following command maps the outside address (209.165.201.15) to an inside address (10.1.1.6):

```
hostname(config)# static (outside,inside) 10.1.1.6 209.165.201.15 netmask 255.255.255.255
```

The following command statically maps an entire subnet:


```
hostname(config)# static (inside,dmz) 10.1.1.0 10.1.2.0 netmask 255.255.255.0
```

Using Static PAT

This section describes how to configure a static port translation. Static PAT lets you translate the real IP address to a mapped IP address, as well as the real port to a mapped port. You can choose to translate the real port to the same port, which lets you translate only specific types of traffic, or you can take it further by translating to a different port.

Figure 17-23 shows a typical static PAT scenario. The translation is always active so both translated and remote hosts can originate connections, and the mapped address and port is statically assigned by the **static** command.

Figure 17-23 Static PAT



For applications that require application inspection for secondary channels (FTP, VoIP, etc.), the security appliance automatically translates the secondary ports.

You cannot use the same real or mapped address in multiple **static** statements between the same two interfaces. Do not use a mapped address in the **static** command that is also defined in a **global** command for the same mapped interface.

For more information about static PAT, see the [“Static PAT” section on page 17-8](#).

**Note**

If you remove a **static** command, existing connections that use the translation are not affected. To remove these connections, enter the **clear local-host** command.

You cannot clear static translations from the translation table with the **clear xlate** command; you must remove the **static** command instead. Only dynamic translations created by the **nat** and **global** commands can be removed with the **clear xlate** command.

To configure static PAT, enter one of the following commands.

- For policy static PAT, enter the following command:

```
hostname(config)# static (real_interface,mapped_interface) {tcp | udp}
{mapped_ip | interface} mapped_port access-list acl_name [dns] [norandomseq]
[[tcp] tcp_max_conns [emb_limit]] [udp udp_max_conns]
```

Create the access list using the **access-list** command (see the [“Adding an Extended Access List” section on page 16-5](#)). The protocol in the access list must match the protocol you set in this command. For example, if you specify **tcp** in the **static** command, then you must specify **tcp** in the access list. Specify the port using the **eq** operator. This access list should include only **permit** ACEs. The source subnet mask used in the access list is also used for the mapped addresses. Policy NAT does not consider the **inactive** or **time-range** keywords; all ACEs are considered to be active for policy NAT configuration.

If you specify a network for translation (for example, 10.1.1.0 255.255.255.0), then the security appliance translates the .0 and .255 addresses. If you want to prevent access to these addresses, be sure to configure an access list to deny access.

See the [“Configuring Dynamic NAT or PAT” section on page 17-23](#) for information about the other options.

- To configure regular static PAT, enter the following command:

```
hostname(config)# static (real_interface,mapped_interface) {tcp | udp} {mapped_ip |
interface} mapped_port real_ip real_port [netmask mask] [dns] [norandomseq] [[tcp]
tcp_max_conns [emb_limit]] [udp udp_max_conns]
```

See the [“Configuring Dynamic NAT or PAT” section on page 17-23](#) for information about the options.

**Note**

When configuring static PAT with FTP, you need to add entries for both TCP ports 20 and 21. You must specify port 20 so that the source port for the active transfer is not modified to another port, which may interfere with other devices that perform NAT on FTP traffic.

For example, for Telnet traffic initiated from hosts on the 10.1.3.0 network to the security appliance outside interface (10.1.2.14), you can redirect the traffic to the inside host at 10.1.1.15 by entering the following commands:

```
hostname(config)# access-list TELNET permit tcp host 10.1.1.15 eq telnet 10.1.3.0
255.255.255.0 eq telnet
hostname(config)# static (inside,outside) tcp 10.1.2.14 telnet access-list TELNET
```

For HTTP traffic initiated from hosts on the 10.1.3.0 network to the security appliance outside interface (10.1.2.14), you can redirect the traffic to the inside host at 10.1.1.15 by entering:

```
hostname(config)# access-list HTTP permit tcp host 10.1.1.15 eq http 10.1.3.0
255.255.255.0 eq http
hostname(config)# static (inside,outside) tcp 10.1.2.14 http access-list HTTP
```

To redirect Telnet traffic from the security appliance outside interface (10.1.2.14) to the inside host at 10.1.1.15, enter the following command:

```
hostname(config)# static (inside,outside) tcp 10.1.2.14 telnet 10.1.1.15 telnet netmask
255.255.255.255
```

If you want to allow the preceding real Telnet server to initiate connections, though, then you need to provide additional translation. For example, to translate all other types of traffic, enter the following commands. The original **static** command provides translation for Telnet to the server, while the **nat** and **global** commands provide PAT for outbound connections from the server.

```
hostname(config)# static (inside,outside) tcp 10.1.2.14 telnet 10.1.1.15 telnet netmask
255.255.255.255
hostname(config)# nat (inside) 1 10.1.1.15 255.255.255.255
hostname(config)# global (outside) 1 10.1.2.14
```

If you also have a separate translation for all inside traffic, and the inside hosts use a different mapped address from the Telnet server, you can still configure traffic initiated from the Telnet server to use the same mapped address as the **static** statement that allows Telnet traffic to the server. You need to create a more exclusive **nat** statement just for the Telnet server. Because **nat** statements are read for the best match, more exclusive **nat** statements are matched before general statements. The following example shows the Telnet **static** statement, the more exclusive **nat** statement for initiated traffic from the Telnet server, and the statement for other inside hosts, which uses a different mapped address.

```
hostname(config)# static (inside,outside) tcp 10.1.2.14 telnet 10.1.1.15 telnet netmask
255.255.255.255
hostname(config)# nat (inside) 1 10.1.1.15 255.255.255.255
hostname(config)# global (outside) 1 10.1.2.14
hostname(config)# nat (inside) 2 10.1.1.0 255.255.255.0
hostname(config)# global (outside) 2 10.1.2.78
```

To translate a well-known port (80) to another port (8080), enter the following command:

```
hostname(config)# static (inside,outside) tcp 10.1.2.45 80 10.1.1.16 8080 netmask
255.255.255.255
```

Bypassing NAT

This section describes how to bypass NAT. You might want to bypass NAT when you enable NAT control. You can bypass NAT using identity NAT, static identity NAT, or NAT exemption. See the [“Bypassing NAT When NAT Control is Enabled” section on page 17-9](#) for more information about these methods. This section includes the following topics:


- [Configuring Identity NAT, page 17-30](#)
- [Configuring Static Identity NAT, page 17-30](#)
- [Configuring NAT Exemption, page 17-32](#)

Configuring Identity NAT

Identity NAT translates the real IP address to the same IP address. Only “translated” hosts can create NAT translations, and responding traffic is allowed back.

Figure 17-24 shows a typical identity NAT scenario.

Figure 17-24 Identity NAT



Note

If you change the NAT configuration, and you do not want to wait for existing translations to time out before the new NAT information is used, you can clear the translation table using the **clear xlate** command. However, clearing the translation table disconnects all current connections that use translations.

To configure identity NAT, enter the following command:

```
hostname(config)# nat (real_interface) 0 real_ip [mask [dns] [outside] [norandomseq]
[[tcp] tcp_max_conns [emb_limit]] [udp udp_max_conns]
```

See the “Configuring Dynamic NAT or PAT” section on page 17-23 for information about the options.

For example, to use identity NAT for the inside 10.1.1.0/24 network, enter the following command:

```
hostname(config)# nat (inside) 0 10.1.1.0 255.255.255.0
```


Configuring Static Identity NAT

Static identity NAT translates the real IP address to the same IP address. The translation is always active, and both “translated” and remote hosts can originate connections. Static identity NAT lets you use regular NAT or policy NAT. Policy NAT lets you identify the real and destination addresses when determining the real addresses to translate (see the “Policy NAT” section on page 17-9 for more

information about policy NAT). For example, you can use policy static identity NAT for an inside address when it accesses the outside interface and the destination is server A, but use a normal translation when accessing the outside server B.

Figure 17-25 shows a typical static identity NAT scenario.

Figure 17-25 Static Identity NAT



Note

If you remove a **static** command, existing connections that use the translation are not affected. To remove these connections, enter the **clear local-host** command.

You cannot clear static translations from the translation table with the **clear xlate** command; you must remove the **static** command instead. Only dynamic translations created by the **nat** and **global** commands can be removed with the **clear xlate** command.

To configure static identity NAT, enter one of the following commands:

- To configure policy static identity NAT, enter the following command:

```
hostname(config)# static (real_interface,mapped_interface) real_ip access-list acl_id
[dns] [norandomseq] [[tcp] tcp_max_conns [emb_limit]] [udp udp_max_conns]
```

Create the access list using the **access-list** command (see the “Adding an Extended Access List” section on page 16-5). This access list should include only **permit** ACEs. Make sure the source address in the access list matches the *real_ip* in this command. Policy NAT does not consider the **inactive** or **time-range** keywords; all ACEs are considered to be active for policy NAT configuration. See the “Policy NAT” section on page 17-9 for more information.

See the “Configuring Dynamic NAT or PAT” section on page 17-23 for information about the other options.

- To configure regular static identity NAT, enter the following command:

```
hostname(config)# static (real_interface,mapped_interface) real_ip real_ip [netmask
mask] [dns] [norandomseq] [[tcp] tcp_max_conns [emb_limit]] [udp udp_max_conns]
```

Specify the same IP address for both *real_ip* arguments.

See the “Configuring Dynamic NAT or PAT” section on page 17-23 for information about the other options.

For example, the following command uses static identity NAT for an inside IP address (10.1.1.3) when accessed by the outside:

```
hostname(config)# static (inside,outside) 10.1.1.3 10.1.1.3 netmask 255.255.255.255
```

The following command uses static identity NAT for an outside address (209.165.201.15) when accessed by the inside:

```
hostname(config)# static (outside,inside) 209.165.201.15 209.165.201.15 netmask
255.255.255.255
```

The following command statically maps an entire subnet:

```
hostname(config)# static (inside,dmz) 10.1.2.0 10.1.2.0 netmask 255.255.255.0
```

The following static identity policy NAT example shows a single real address that uses identity NAT when accessing one destination address, and a translation when accessing another:


```
hostname(config)# access-list NET1 permit ip host 10.1.2.27 209.165.201.0 255.255.255.224
hostname(config)# access-list NET2 permit ip host 10.1.2.27 209.165.200.224
255.255.255.224
hostname(config)# static (inside,outside) 10.1.2.27 access-list NET1
hostname(config)# static (inside,outside) 209.165.202.130 access-list NET2
```

Configuring NAT Exemption

NAT exemption exempts addresses from translation and allows both real and remote hosts to originate connections. NAT exemption lets you specify the real and destination addresses when determining the real traffic to exempt (similar to policy NAT), so you have greater control using NAT exemption than identity NAT. However unlike policy NAT, NAT exemption does not consider the ports in the access list. Use static identity NAT to consider ports in the access list.

Figure 17-26 shows a typical NAT exemption scenario.

Figure 17-26 NAT Exemption



Note

If you remove a NAT exemption configuration, existing connections that use NAT exemption are not affected. To remove these connections, enter the **clear local-host** command.

To configure NAT exemption, enter the following command:

```
hostname(config)# nat (real_interface) 0 access-list acl_name [outside] [norandomseq]
[[tcp] tcp_max_conns [emb_limit]] [udp udp_max_conns]
```

Create the access list using the **access-list** command (see the [“Adding an Extended Access List”](#) section on page 16-5). This access list can include both **permit** ACEs and **deny** ACEs. Do not specify the real and destination ports in the access list; NAT exemption does not consider the ports. NAT exemption considers the **inactive** and **time-range** keywords, but it does not support ACL with all **inactive** and **time-range** ACEs.

See the [“Configuring Dynamic NAT or PAT” section on page 17-23](#) for information about the other options.

By default, this command exempts traffic from inside to outside. If you want traffic from outside to inside to bypass NAT, then add an additional **nat** command and enter **outside** to identify the NAT instance as outside NAT. You might want to use outside NAT exemption if you configure dynamic NAT for the outside interface and want to exempt other traffic.

For example, to exempt an inside network when accessing any destination address, enter the following command:

```
hostname(config)# access-list EXEMPT permit ip 10.1.2.0 255.255.255.0 any
hostname(config)# nat (inside) 0 access-list EXEMPT
```

To use dynamic outside NAT for a DMZ network, and exempt another DMZ network, enter the following command:

```
hostname(config)# nat (dmz) 1 10.1.2.0 255.255.255.0 outside dns
hostname(config)# global (inside) 1 10.1.1.45
hostname(config)# access-list EXEMPT permit ip 10.1.3.0 255.255.255.0 any
hostname(config)# nat (dmz) 0 access-list EXEMPT
```

To exempt an inside address when accessing two different destination addresses, enter the following commands:

```
hostname(config)# access-list NET1 permit ip 10.1.2.0 255.255.255.0 209.165.201.0
255.255.255.224
hostname(config)# access-list NET1 permit ip 10.1.2.0 255.255.255.0 209.165.200.224
255.255.255.224
hostname(config)# nat (inside) 0 access-list NET1
```

NAT Examples


This section describes typical scenarios that use NAT solutions, and includes the following topics:

- [Overlapping Networks, page 17-34](#)
- [Redirecting Ports, page 17-35](#)

Overlapping Networks

In Figure 17-27, the security appliance connects two private networks with overlapping address ranges.

Figure 17-27 Using Outside NAT with Overlapping Networks



Two networks use an overlapping address space (192.168.100.0/24), but hosts on each network must communicate (as allowed by access lists). Without NAT, when a host on the inside network tries to access a host on the overlapping DMZ network, the packet never makes it past the security appliance, which sees the packet as having a destination address on the inside network. Moreover, if the destination address is being used by another host on the inside network, that host receives the packet.

To solve this problem, use NAT to provide non-overlapping addresses. If you want to allow access in both directions, use static NAT for both networks. If you only want to allow the inside interface to access hosts on the DMZ, then you can use dynamic NAT for the inside addresses, and static NAT for the DMZ addresses you want to access. This example shows static NAT.

To configure static NAT for these two interfaces, perform the following steps. The 10.1.1.0/24 network on the DMZ is not translated.

-
- Step 1** Translate 192.168.100.0/24 on the inside to 10.1.2.0 /24 when it accesses the DMZ by entering the following command:
- ```
hostname(config)# static (inside,dmz) 10.1.2.0 192.168.100.0 netmask 255.255.255.0
```
- Step 2** Translate the 192.168.100.0/24 network on the DMZ to 10.1.3.0/24 when it accesses the inside by entering the following command:
- ```
hostname(config)# static (dmz,inside) 10.1.3.0 192.168.100.0 netmask 255.255.255.0
```
- Step 3** Configure the following static routes so that traffic to the dmz network can be routed correctly by the security appliance:
- ```
hostname(config)# route dmz 192.168.100.128 255.255.255.128 10.1.1.2 1
hostname(config)# route dmz 192.168.100.0 255.255.255.128 10.1.1.2 1
```

The security appliance already has a connected route for the inside network. These static routes allow the security appliance to send traffic for the 192.168.100.0/24 network out the DMZ interface to the gateway router at 10.1.1.2. (You need to split the network into two because you cannot create a static route with the exact same network as a connected route.) Alternatively, you could use a more broad route for the DMZ traffic, such as a default route.


If host 192.168.100.2 on the DMZ network wants to initiate a connection to host 192.168.100.2 on the inside network, the following events occur:

1. The DMZ host 192.168.100.2 sends the packet to IP address 10.1.2.2.
2. When the security appliance receives this packet, the security appliance translates the source address from 192.168.100.2 to 10.1.3.2.
3. Then the security appliance translates the destination address from 10.1.2.2 to 192.168.100.2, and the packet is forwarded.

## Redirecting Ports

Figure 17-28 illustrates a typical network scenario in which the port redirection feature might be useful.

**Figure 17-28 Port Redirection Using Static PAT**



In the configuration described in this section, port redirection occurs for hosts on external networks as follows:

- Telnet requests to IP address 209.165.201.5 are redirected to 10.1.1.6.
- FTP requests to IP address 209.165.201.5 are redirected to 10.1.1.3.
- HTTP request to security appliance outside IP address 209.165.201.25 are redirected to 10.1.1.5.
- HTTP port 8080 requests to PAT address 209.165.201.15 are redirected to 10.1.1.7 port 80.

To implement this scenario, perform the following steps:

---

**Step 1** Configure PAT for the inside network by entering the following commands:

```
hostname(config)# nat (inside) 1 0.0.0.0 0.0.0.0 0 0
hostname(config)# global (outside) 1 209.165.201.15
```

**Step 2** Redirect Telnet requests for 209.165.201.5 to 10.1.1.6 by entering the following command:

```
hostname(config)# static (inside,outside) tcp 209.165.201.5 telnet 10.1.1.6 telnet netmask
255.255.255.255
```

**Step 3** Redirect FTP requests for IP address 209.165.201.5 to 10.1.1.3 by entering the following command:

```
hostname(config)# static (inside,outside) tcp 209.165.201.5 ftp 10.1.1.3 ftp netmask
255.255.255.255
```

**Step 4** Redirect HTTP requests for the security appliance outside interface address to 10.1.1.5 by entering the following command:

```
hostname(config)# static (inside,outside) tcp interface www 10.1.1.5 www netmask
255.255.255.255
```

**Step 5** Redirect HTTP requests on port 8080 for PAT address 209.165.201.15 to 10.1.1.7 port 80 by entering the following command:

```
hostname(config)# static (inside,outside) tcp 209.165.201.15 8080 10.1.1.7 www netmask
255.255.255.255
```

---



# CHAPTER 18

## Permitting or Denying Network Access

This chapter describes how to control network access through the security appliance using access lists. To create an extended access lists or an EtherType access list, see [Chapter 16, “Identifying Traffic with Access Lists.”](#)



### Note

You use ACLs to control network access in both routed and transparent firewall modes. In transparent mode, you can use both extended ACLs (for Layer 3 traffic) and EtherType ACLs (for Layer 2 traffic).

To access the security appliance interface for management access, you do not also need an access list allowing the host IP address. You only need to configure management access according to [Chapter 40, “Managing System Access.”](#)

This chapter includes the following sections:

- [Inbound and Outbound Access List Overview, page 18-1](#)
- [Applying an Access List to an Interface, page 18-2](#)

## Inbound and Outbound Access List Overview

By default, all traffic from a higher-security interface to a lower-security interface is allowed. Access lists let you either allow traffic from lower-security interfaces, or restrict traffic from higher-security interfaces.

The security appliance supports two types of access lists:

- Inbound—Inbound access lists apply to traffic as it enters an interface.
- Outbound—Outbound access lists apply to traffic as it exits an interface.




### Note

“Inbound” and “outbound” refer to the application of an access list on an interface, either to traffic entering the security appliance on an interface or traffic exiting the security appliance on an interface. These terms do not refer to the movement of traffic from a lower security interface to a higher security interface, commonly known as inbound, or from a higher to lower interface, commonly known as outbound.

An outbound access list is useful, for example, if you want to allow only certain hosts on the inside networks to access a web server on the outside network. Rather than creating multiple inbound access lists to restrict access, you can create a single outbound access list that allows only the specified hosts

(see Figure 18-1). See the “IP Addresses Used for Access Lists When You Use NAT” section on page 16-3 for information about NAT and IP addresses. The outbound access list prevents any other hosts from reaching the outside network.

**Figure 18-1 Outbound Access List**



See the following commands for this example:

```
hostname(config)# access-list OUTSIDE extended permit tcp host 209.165.201.4
host 209.165.200.225 eq www
hostname(config)# access-list OUTSIDE extended permit tcp host 209.165.201.6
host 209.165.200.225 eq www
hostname(config)# access-list OUTSIDE extended permit tcp host 209.165.201.8
host 209.165.200.225 eq www
hostname(config)# access-group OUTSIDE out interface outside
```

## Applying an Access List to an Interface

To apply an extended access list to the inbound or outbound direction of an interface, enter the following command:

```
hostname(config)# access-group access_list_name {in | out} interface interface_name
[per-user-override]
```

You can apply one access list of each type (extended and EtherType) to both directions of the interface. See the “Inbound and Outbound Access List Overview” section on page 18-1 for more information about access list directions.



The **per-user-override** keyword allows dynamic access lists that are downloaded for user authorization to override the access list assigned to the interface. For example, if the interface access list denies all traffic from 10.0.0.0, but the dynamic access list permits all traffic from 10.0.0.0, then the dynamic access list overrides the interface access list for that user. See the “Configuring RADIUS Authorization” section for more information about per-user access lists. The **per-user-override** keyword is only available for inbound access lists.

For connectionless protocols, you need to apply the access list to the source and destination interfaces if you want traffic to pass in both directions.

The following example illustrates the commands required to enable access to an inside web server with the IP address 209.165.201.12 (this IP address is the address visible on the outside interface after NAT):

```
hostname(config)# access-list ACL_OUT extended permit tcp any host 209.165.201.12 eq www
hostname(config)# access-group ACL_OUT in interface outside
```

You also need to configure NAT for the web server.

The following access lists allow *any* hosts to communicate between the **inside** and **hr** networks, but only specific hosts (209.168.200.3 and 209.168.200.4) to access the outside network, as shown in the last line below:

```
hostname(config)# access-list ANY extended permit ip any any
hostname(config)# access-list OUT extended permit ip host 209.168.200.3 any
hostname(config)# access-list OUT extended permit ip host 209.168.200.4 any

hostname(config)# access-group ANY in interface inside
hostname(config)# access-group ANY in interface hr
hostname(config)# access-group OUT out interface outside
```

For example, the following sample access list allows common EtherTypes originating on the inside interface:

```
hostname(config)# access-list ETHER ethertype permit ipx
hostname(config)# access-list ETHER ethertype permit bpdu
hostname(config)# access-list ETHER ethertype permit mpls-unicast
hostname(config)# access-group ETHER in interface inside
```

The following access list allows some EtherTypes through the security appliance, but denies all others:

```
hostname(config)# access-list ETHER ethertype permit 0x1234
hostname(config)# access-list ETHER ethertype permit bpdu
hostname(config)# access-list ETHER ethertype permit mpls-unicast
hostname(config)# access-group ETHER in interface inside
hostname(config)# access-group ETHER in interface outside
```

The following access list denies traffic with EtherType 0x1256 but allows all others on both interfaces:

```
hostname(config)# access-list nonIP ethertype deny 1256
hostname(config)# access-list nonIP ethertype permit any
hostname(config)# access-group ETHER in interface inside
hostname(config)# access-group ETHER in interface outside
```





# CHAPTER 19

## Applying AAA for Network Access

---

This chapter describes how to enable AAA (pronounced “triple A”) for network access.

For information about AAA for management access, see the [“Configuring AAA for System Administrators”](#) section on page 40-5.

This chapter contains the following sections:

- [AAA Performance, page 19-1](#)
- [Configuring Authentication for Network Access, page 19-1](#)
- [Configuring Authorization for Network Access, page 19-6](#)
- [Configuring Accounting for Network Access, page 19-13](#)
- [Using MAC Addresses to Exempt Traffic from Authentication and Authorization, page 19-14](#)

### AAA Performance

The security appliance uses “cut-through proxy” to significantly improve performance compared to a traditional proxy server. The performance of a traditional proxy server suffers because it analyzes every packet at the application layer of the OSI model. The security appliance cut-through proxy challenges a user initially at the application layer and then authenticates against standard AAA servers or the local database. After the security appliance authenticates the user, it shifts the session flow, and all traffic flows directly and quickly between the source and destination while maintaining session state information.

### Configuring Authentication for Network Access

This section includes the following topics:

- [Authentication Overview, page 19-2](#)
- [Enabling Network Access Authentication, page 19-3](#)
- [Enabling Secure Authentication of Web Clients, page 19-5](#)
- [Authenticating Directly with the Security Appliance, page 19-6](#)

## Authentication Overview

The security appliance lets you configure network access authentication using AAA servers. This section includes the following topics:

- [One-Time Authentication, page 19-2](#)
- [Applications Required to Receive an Authentication Challenge, page 19-2](#)
- [Security Appliance Authentication Prompts, page 19-2](#)
- [Static PAT and HTTP, page 19-3](#)
- [Enabling Network Access Authentication, page 19-3](#)

### One-Time Authentication

A user at a given IP address only needs to authenticate one time for all rules and types, until the authentication session expires. (See the **timeout uauth** command in the *Cisco Security Appliance Command Reference* for timeout values.) For example, if you configure the security appliance to authenticate Telnet and FTP, and a user first successfully authenticates for Telnet, then as long as the authentication session exists, the user does not also have to authenticate for FTP.

### Applications Required to Receive an Authentication Challenge

Although you can configure the security appliance to require authentication for network access to any protocol or service, users can authenticate directly with HTTP, HTTPS, Telnet, or FTP only. A user must first authenticate with one of these services before the security appliance allows other traffic requiring authentication.

The authentication ports that the security appliance supports for AAA are fixed:

- Port 21 for FTP
- Port 23 for Telnet
- Port 80 for HTTP
- Port 443 for HTTPS

### Security Appliance Authentication Prompts

For Telnet and FTP, the security appliance generates an authentication prompt.

For HTTP, the security appliance uses basic HTTP authentication by default, and provides an authentication prompt. You can optionally configure the security appliance to redirect users to an internal web page where they can enter their username and password (configured with the **aaa authentication listener** command).

For HTTPS, the security appliance generates a custom login screen. You can optionally configure the security appliance to redirect users to an internal web page where they can enter their username and password (configured with the **aaa authentication listener** command).

Redirection is an improvement over the basic method because it provides an improved user experience when authenticating, and an identical user experience for HTTP and HTTPS in both Easy VPN and firewall modes. It also supports authenticating directly with the security appliance.

You might want to continue to use basic HTTP authentication if: you do not want the security appliance to open listening ports; if you use NAT on a router and you do not want to create a translation rule for the web page served by the security appliance; basic HTTP authentication might work better with your network. For example non-browser applications, like when a URL is embedded in email, might be more compatible with basic authentication.

After you authenticate correctly, the security appliance redirects you to your original destination. If the destination server also has its own authentication, the user enters another username and password. If you use basic HTTP authentication and need to enter another username and password for the destination server, then you need to configure the **virtual http** command.

**Note**

If you use HTTP authentication without using the **aaa authentication secure-http-client** command, the username and password are sent from the client to the security appliance in clear text. We recommend that you use the **aaa authentication secure-http-client** command whenever you enable HTTP authentication. For more information about the **aaa authentication secure-http-client** command, see the [“Enabling Secure Authentication of Web Clients” section on page 19-5](#).

For FTP, a user has the option of entering the security appliance username followed by an at sign (@) and then the FTP username (name1@name2). For the password, the user enters the security appliance password followed by an at sign (@) and then the FTP password (password1@password2). For example, enter the following text.

```
name> jamiec@jchrichton
password> letmein@he110
```

This feature is useful when you have cascaded firewalls that require multiple logins. You can separate several names and passwords by multiple at signs (@).

## Static PAT and HTTP

For HTTP authentication, the security appliance checks real ports when static PAT is configured. If it detects traffic destined for real port 80, regardless of the mapped port, the security appliance intercepts the HTTP connection and enforces authentication.

For example, assume that outside TCP port 889 is translated to port 80 (www) and that any relevant access lists permit the traffic:

```
static (inside,outside) tcp 10.48.66.155 889 192.168.123.10 www netmask 255.255.255.255
```

Then when users try to access 10.48.66.155 on port 889, the security appliance intercepts the traffic and enforces HTTP authentication. Users see the HTTP authentication page in their web browsers before the security appliance allows HTTP connection to complete.

If the local port is different than port 80, as in the following example:

```
static (inside,outside) tcp 10.48.66.155 889 192.168.123.10 111 netmask 255.255.255.255
```

Then users do not see the authentication page. Instead, the security appliance sends to the web browser an error message indicating that the user must be authenticated prior using the requested service.

## Enabling Network Access Authentication

To enable network access authentication, perform the following steps:

- Step 1** Using the **aaa-server** command, identify your AAA servers. If you have already identified your AAA servers, continue to the next step.

For more information about identifying AAA servers, see the [“Identifying AAA Server Groups and Servers” section on page 13-12](#).

- Step 2** Using the **access-list** command, create an access list that identifies the source addresses and destination addresses of traffic you want to authenticate. For steps, see the [“Adding an Extended Access List” section on page 16-5](#).

The **permit** ACEs mark matching traffic for authentication, while **deny** entries exclude matching traffic from authentication. Be sure to include the destination ports for either HTTP, HTTPS, Telnet, or FTP in the access list because the user must authenticate with one of these services before other services are allowed through the security appliance.

- Step 3** To configure authentication, enter the following command:

```
hostname(config)# aaa authentication match acl_name interface_name server_group
```

Where *acl\_name* is the name of the access list you created in [Step 2](#), *interface\_name* is the name of the interface as specified with the **nameif** command, and *server\_group* is the AAA server group you created in [Step 1](#).



**Note**

You can alternatively use the **aaa authentication include** command (which identifies traffic within the command). However, you cannot use both methods in the same configuration. See the *Cisco Security Appliance Command Reference* for more information.

- Step 4** (Optional) To enable the redirection method of authentication for HTTP or HTTPS connections, enter the following command:

```
hostname(config)# aaa authentication listener http[s] interface_name [port portnum]
redirect
```

where the *interface\_name* argument is the interface on which you want to enable listening ports.

The **port portnum** argument specifies the port number that the security appliance listens on; the defaults are 80 (HTTP) and 443 (HTTPS).

Enter this command separately for HTTP and for HTTPS.

- Step 5** (Optional) If you are using the local database for network access authentication and you want to limit the number of consecutive failed login attempts that the security appliance allows any given user account, use the following command:

```
hostname(config)# aaa local authentication attempts max-fail number
```

Where *number* is between 1 and 16.

For example:

```
hostname(config)# aaa local authentication attempts max-fail 7
```



**Tip**

To clear the lockout status of a specific user or all users, use the **clear aaa local user lockout** command.

For example, the following commands authenticate all inside HTTP traffic and SMTP traffic:

```
hostname(config)# aaa-server AuthOutbound protocol tacacs+
```

```
hostname(config-aaa-server-group)# exit
hostname(config)# aaa-server AuthOutbound (inside) host 10.1.1.1
hostname(config-aaa-server-host)# key TACPlusUauthKey
hostname(config-aaa-server-host)# exit
hostname(config)# access-list MAIL_AUTH extended permit tcp any any eq smtp
hostname(config)# access-list MAIL_AUTH extended permit tcp any any eq www
hostname(config)# aaa authentication match MAIL_AUTH inside AuthOutbound
hostname(config)# aaa authentication listener http inside redirect
```

The following commands authenticate Telnet traffic from the outside interface to a particular server (209.165.201.5):

```
hostname(config)# aaa-server AuthInbound protocol tacacs+
hostname(config-aaa-server-group)# exit
hostname(config)# aaa-server AuthInbound (inside) host 10.1.1.1
hostname(config-aaa-server-host)# key TACPlusUauthKey
hostname(config-aaa-server-host)# exit
hostname(config)# access-list TELNET_AUTH extended permit tcp any host 209.165.201.5 eq telnet
hostname(config)# aaa authentication match TELNET_AUTH outside AuthInbound
```

## Enabling Secure Authentication of Web Clients

The security appliance provides a method of securing HTTP authentication. Without securing HTTP authentication, usernames and passwords from the client to the security appliance would be passed as clear text. By using the **aaa authentication secure-http-client** command, you enable the exchange of usernames and passwords between a web client and the security appliance with HTTPS.

After enabling this feature, when a user requires authentication when using HTTP, the security appliance redirects the HTTP user to an HTTPS prompt. After you authenticate correctly, the security appliance redirects you to the original HTTP URL.

To enable secure authentication of web clients, enter the following command:

```
hostname(config)# aaa authentication secure-http-client
```

Secured web-client authentication has the following limitations:

- A maximum of 16 concurrent HTTPS authentication sessions are allowed. If all 16 HTTPS authentication processes are running, a new connection requiring authentication will not succeed.
- When **uauth timeout 0** is configured (the **uauth timeout** is set to 0), HTTPS authentication might not work. If a browser initiates multiple TCP connections to load a web page after HTTPS authentication, the first connection is let through, but the subsequent connections trigger authentication. As a result, users are continuously presented with an authentication page, even if the correct username and password are entered each time. To work around this, set the **uauth timeout** to 1 second with the **timeout uauth 0:0:1** command. However, this workaround opens a 1-second window of opportunity that might allow non-authenticated users to go through the firewall if they are coming from the same source IP address.
- Because HTTPS authentication occurs on the SSL port 443, users must not configure an **access-list** command statement to block traffic from the HTTP client to HTTP server on port 443. Furthermore, if static PAT is configured for web traffic on port 80, it must also be configured for the SSL port. In the following example, the first line configures static PAT for web traffic and the second line must be added to support the HTTPS authentication configuration.

```
static (inside,outside) tcp 10.132.16.200 www 10.130.16.10 www
static (inside,outside) tcp 10.132.16.200 443 10.130.16.10 443
```

## Authenticating Directly with the Security Appliance

If you do not want to allow HTTP, HTTPS, Telnet, or FTP through the security appliance but want to authenticate other types of traffic, you can authenticate with the security appliance directly using HTTP, HTTPS, or Telnet.

This section includes the following topics:

- [Enabling Direct Authentication Using HTTP and HTTPS, page 19-6](#)
- [Enabling Direct Authentication Using Telnet, page 19-6](#)

### Enabling Direct Authentication Using HTTP and HTTPS

If you enabled the redirect method of HTTP and HTTPS authentication in the [“Enabling Network Access Authentication” section on page 19-3](#), then you also automatically enabled direct authentication. If you want to continue to use basic HTTP authentication, but want to enable direct authentication for HTTP and HTTPS, then enter the following command:

```
hostname(config)# aaa authentication listener http[s] interface_name [port portnum]
```

where the *interface\_name* argument is the interface on which you want to enable direct authentication.

The **port portnum** argument specifies the port number that the security appliance listens on; the defaults are 80 (HTTP) and 443 (HTTPS).

Enter this command separately for HTTP and for HTTPS.

You can authenticate directly with the security appliance at the following URLs when you enable AAA for the interface:

```
http://interface_ip[:port]/netaccess/connstatus.html
https://interface_ip[:port]/netaccess/connstatus.html
```

### Enabling Direct Authentication Using Telnet

To enable direct authentication with Telnet, configure a virtual Telnet server. With virtual Telnet, the user Telnets to a given IP address configured on the security appliance, and the security appliance provides a Telnet prompt. To configure a virtual Telnet server, enter the following command:

```
hostname(config)# virtual telnet ip_address
```

where the *ip\_address* argument sets the IP address for the virtual Telnet server. Make sure this address is an unused address that is routed to the security appliance. For example, if you perform NAT for inside addresses when they access the outside, and you want to provide outside access to the virtual Telnet server, you can use one of the global NAT addresses for the virtual Telnet server address.

## Configuring Authorization for Network Access

After a user authenticates for a given connection, the security appliance can use authorization to further control traffic from the user.

This section includes the following topics:

- [Configuring TACACS+ Authorization, page 19-7](#)
- [Configuring RADIUS Authorization, page 19-8](#)



## Configuring TACACS+ Authorization

You can configure the security appliance to perform network access authorization with TACACS+. You identify the traffic to be authorized by specifying access lists that authorization rules must match. Alternatively, you can identify the traffic directly in authorization rules themselves.



**Tip**

Using access lists to identify traffic to be authorized can greatly reduced the number of authorization commands you must enter. This is because each authorization rule you enter can specify only one source and destination subnet and service, whereas an access list can include many entries.

Authentication and authorization statements are independent; however, any unauthenticated traffic matched by an authorization statement will be denied. For authorization to succeed, a user must first authenticate with the security appliance. Because a user at a given IP address only needs to authenticate one time for all rules and types, if the authentication session hasn't expired, authorization can occur even if the traffic is matched by an authentication statement.

After a user authenticates, the security appliance checks the authorization rules for matching traffic. If the traffic matches the authorization statement, the security appliance sends the username to the TACACS+ server. The TACACS+ server responds to the security appliance with a permit or a deny for that traffic, based on the user profile. The security appliance enforces the authorization rule in the response.

See the documentation for your TACACS+ server for information about configuring network access authorizations for a user.

To configure TACACS+ authorization, perform the following steps:

**Step 1** Enable authentication. For more information, see the [“Enabling Network Access Authentication” section on page 19-3](#). If you have already enabled authentication, continue to the next step.

**Step 2** Using the **access-list** command, create an access list that identifies the source addresses and destination addresses of traffic you want to authorize. For steps, see the [“Adding an Extended Access List” section on page 16-5](#).

The **permit** ACEs mark matching traffic for authorization, while **deny** entries exclude matching traffic from authorization. The access list you use for authorization matching should contain rules that are equal to or a subset of the rules in the access list used for authentication matching.



**Note**

If you have configured authentication and want to authorize all the traffic being authenticated, you can use the same access list you created for use with the **aaa authentication match** command.

**Step 3** To enable authorization, enter the following command:

```
hostname(config)# aaa authorization match acl_name interface_name server_group
```

where *acl\_name* is the name of the access list you created in [Step 2](#), *interface\_name* is the name of the interface as specified with the **nameif** command or by default, and *server\_group* is the AAA server group you created when you enabled authentication.

**Note**

Alternatively, you can use the **aaa authorization include** command (which identifies traffic within the command) but you cannot use both methods in the same configuration. See the *Cisco Security Appliance Command Reference* for more information.

The following commands authenticate and authorize inside Telnet traffic. Telnet traffic to servers other than 209.165.201.5 can be authenticated alone, but traffic to 209.165.201.5 requires authorization.

```
hostname(config)# access-list TELNET_AUTH extended permit tcp any any eq telnet
hostname(config)# access-list SERVER_AUTH extended permit tcp any host 209.165.201.5 eq telnet
hostname(config)# aaa-server AuthOutbound protocol tacacs+
hostname(config-aaa-server-group)# exit
hostname(config)# aaa-server AuthOutbound (inside) host 10.1.1.1
hostname(config-aaa-server-host)# key TACPlusUauthKey
hostname(config-aaa-server-host)# exit
hostname(config)# aaa authentication match TELNET_AUTH inside AuthOutbound
hostname(config)# aaa authorization match SERVER_AUTH inside AuthOutbound
```

## Configuring RADIUS Authorization

When authentication succeeds, the RADIUS protocol returns user authorizations in the access-accept message sent by a RADIUS server. For more information about configuring authentication, see the [“Configuring Authentication for Network Access” section on page 19-1](#).

When you configure the security appliance to authenticate users for network access, you are also implicitly enabling RADIUS authorizations; therefore, this section contains no information about configuring RADIUS authorization on the security appliance. It does provide information about how the security appliance handles access list information received from RADIUS servers.

You can configure a RADIUS server to download an access list to the security appliance or an access list name at the time of authentication. The user is authorized to do only what is permitted in the user-specific access list.

**Note**

If you have used the **access-group** command to apply access lists to interfaces, be aware of the following effects of the **per-user-override** keyword on authorization by user-specific access lists:

- Without the **per-user-override** keyword, traffic for a user session must be permitted by both the interface access list and the user-specific access list.
- With the **per-user-override** keyword, the user-specific access list determines what is permitted.

For more information, see the **access-group** command entry in the *Cisco Security Appliance Command Reference*.

This section includes the following topics:

- [Configuring a RADIUS Server to Send Downloadable Access Control Lists, page 19-9](#)
- [Configuring a RADIUS Server to Download Per-User Access Control List Names, page 19-12](#)

## Configuring a RADIUS Server to Send Downloadable Access Control Lists

This section describes how to configure Cisco Secure ACS or a third-party RADIUS server, and includes the following topics:

- [About the Downloadable Access List Feature and Cisco Secure ACS, page 19-9](#)
- [Configuring Cisco Secure ACS for Downloadable Access Lists, page 19-10](#)
- [Configuring Any RADIUS Server for Downloadable Access Lists, page 19-11](#)
- [Converting Wildcard Netmask Expressions in Downloadable Access Lists, page 19-12](#)

### About the Downloadable Access List Feature and Cisco Secure ACS

Downloadable access lists is the most scalable means of using Cisco Secure ACS to provide the appropriate access lists for each user. It provides the following capabilities:

- Unlimited access list size—Downloadable access lists are sent using as many RADIUS packets as required to transport the full access list from Cisco Secure ACS to the security appliance.
- Simplified and centralized management of access lists—Downloadable access lists enable you to write a set of access lists once and apply it to many user or group profiles and distribute it to many security appliances.

This approach is most useful when you have very large access list sets that you want to apply to more than one Cisco Secure ACS user or group; however, its ability to simplify Cisco Secure ACS user and group management makes it useful for access lists of any size.

The security appliance receives downloadable access lists from Cisco Secure ACS using the following process:

1. The security appliance sends a RADIUS authentication request packet for the user session.
2. If Cisco Secure ACS successfully authenticates the user, Cisco Secure ACS returns a RADIUS access-accept message that contains the internal name of the applicable downloadable access list. The Cisco IOS cisco-av-pair RADIUS VSA (vendor 9, attribute 1) contains the following attribute-value pair to identify the downloadable access list set:

```
ACS:CiscoSecure-Defined-ACL=acl-set-name
```

where *acl-set-name* is the internal name of the downloadable access list, which is a combination of the name assigned to the access list by the Cisco Secure ACS administrator and the date and time that the access list was last modified.

3. The security appliance examines the name of the downloadable access list and determines if it has previously received the named downloadable access list.
  - If the security appliance has previously received the named downloadable access list, communication with Cisco Secure ACS is complete and the security appliance applies the access list to the user session. Because the name of the downloadable access list includes the date and time it was last modified, matching the name sent by Cisco Secure ACS to the name of an access list previously downloaded means that the security appliance has the most recent version of the downloadable access list.
  - If the security appliance has not previously received the named downloadable access list, it may have an out-of-date version of the access list or it may not have downloaded any version of the access list. In either case, the security appliance issues a RADIUS authentication request using the downloadable access list name as the username in the RADIUS request and a null password attribute. In a cisco-av-pair RADIUS VSA, the request also includes the following attribute-value pairs:

```
AAA:service=ip-admission
AAA:event=acl-download
```

In addition, the security appliance signs the request with the Message-Authenticator attribute (IETF RADIUS attribute 80).

4. Upon receipt of a RADIUS authentication request that has a username attribute containing the name of a downloadable access list, Cisco Secure ACS authenticates the request by checking the Message-Authenticator attribute. If the Message-Authenticator attribute is missing or incorrect, Cisco Secure ACS ignores the request. The presence of the Message-Authenticator attribute prevents malicious use of a downloadable access list name to gain unauthorized network access. The Message-Authenticator attribute and its use are defined in RFC 2869, RADIUS Extensions, available at <http://www.ietf.org>.
5. If the access list required is less than approximately 4 KB in length, Cisco Secure ACS responds with an access-accept message containing the access list. The largest access list that can fit in a single access-accept message is slightly less than 4 KB because some of the message must be other required attributes.

Cisco Secure ACS sends the downloadable access list in a cisco-av-pair RADIUS VSA. The access list is formatted as a series of attribute-value pairs that each contain an ACE and are numbered serially:

```
ip:inacl#1=ACE-1
ip:inacl#2=ACE-2
.
.
ip:inacl#n=ACE-n
```

An example of an attribute-value pair follows:

```
ip:inacl#1=permit tcp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
```

6. If the access list required is more than approximately 4 KB in length, Cisco Secure ACS responds with an access-challenge message that contains a portion of the access list, formatted as described above, and an State attribute (IETF RADIUS attribute 24), which contains control data used by Cisco Secure ACS to track the progress of the download. Cisco Secure ACS fits as many complete attribute-value pairs into the cisco-av-pair RADIUS VSA as it can without exceeding the maximum RADIUS message size.

The security appliance stores the portion of the access list received and responds with another access-request message containing the same attributes as the first request for the downloadable access list plus a copy of the State attribute received in the access-challenge message.

This repeats until Cisco Secure ACS sends the last of the access list in an access-accept message.

## Configuring Cisco Secure ACS for Downloadable Access Lists

You can configure downloadable access lists on Cisco Secure ACS as a shared profile component and then assign the access list to a group or to an individual user.

The access list definition consists of one or more security appliance commands that are similar to the extended **access-list** command (see the “[Adding an Extended Access List](#)” section on page 16-5), except without the following prefix:

```
access-list acl_name extended
```

The following example is a downloadable access list definition on Cisco Secure ACS version 3.3:

```
+-----+
```

```

| Shared profile Components
|
| Downloadable IP ACLs Content
|
| Name: acs_ten_acl
|
| ACL Definitions
|
| permit tcp any host 10.0.0.254
| permit udp any host 10.0.0.254
| permit icmp any host 10.0.0.254
| permit tcp any host 10.0.0.253
| permit udp any host 10.0.0.253
| permit icmp any host 10.0.0.253
| permit tcp any host 10.0.0.252
| permit udp any host 10.0.0.252
| permit icmp any host 10.0.0.252
| permit ip any any
+-----+

```

For more information about creating downloadable access lists and associating them with users, see the user guide for your version of Cisco Secure ACS.

On the security appliance, the downloaded access list has the following name:

```
#ACSACL#-ip-acl_name-number
```

The *acl\_name* argument is the name that is defined on Cisco Secure ACS (*acs\_ten\_acl* in the preceding example), and *number* is a unique version ID generated by Cisco Secure ACS.

The downloaded access list on the security appliance consists of the following lines:

```

access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit tcp any host 10.0.0.254
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit udp any host 10.0.0.254
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit icmp any host 10.0.0.254
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit tcp any host 10.0.0.253
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit udp any host 10.0.0.253
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit icmp any host 10.0.0.253
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit tcp any host 10.0.0.252
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit udp any host 10.0.0.252
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit icmp any host 10.0.0.252
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit ip any any

```

## Configuring Any RADIUS Server for Downloadable Access Lists

You can configure any RADIUS server that supports Cisco IOS RADIUS VSAs to send user-specific access lists to the security appliance in a Cisco IOS RADIUS *cisco-av-pair* VSA (vendor 9, attribute 1).

In the *cisco-av-pair* VSA, configure one or more ACEs that are similar to the **access-list extended** command (see the [“Adding an Extended Access List”](#) section on page 16-5), except that you replace the following command prefix:

```
access-list acl_name extended
```

with the following text:

```
ip:inacl#nnn=
```

The *nnn* argument is a number in the range from 0 to 999999999 that identifies the order of the command statement to be configured on the security appliance. If this parameter is omitted, the sequence value is 0, and the order of the ACEs inside the *cisco-av-pair* RADIUS VSA is used.

The following example is an access list definition as it should be configured for a cisco-av-pair VSA on a RADIUS server:

```
ip:inacl#1=permit tcp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
ip:inacl#99=deny tcp any any
ip:inacl#2=permit udp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
ip:inacl#100=deny udp any any
ip:inacl#3=permit icmp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
```

For information about making unique per user the access lists that are sent in the cisco-av-pair attribute, see the documentation for your RADIUS server.

On the security appliance, the downloaded access list name has the following format:

```
AAA-user-username
```

The *username* argument is the name of the user that is being authenticated.

The downloaded access list on the security appliance consists of the following lines. Notice the order based on the numbers identified on the RADIUS server.

```
access-list AAA-user-bcham34-79AD4A08 permit tcp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
access-list AAA-user-bcham34-79AD4A08 permit udp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
access-list AAA-user-bcham34-79AD4A08 permit icmp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
access-list AAA-user-bcham34-79AD4A08 deny tcp any any
access-list AAA-user-bcham34-79AD4A08 deny udp any any
```

Downloaded access lists have two spaces between the word “access-list” and the name. These spaces serve to differentiate a downloaded access list from a local access list. In this example, “79AD4A08” is a hash value generated by the security appliance to help determine when access list definitions have changed on the RADIUS server.

## Converting Wildcard Netmask Expressions in Downloadable Access Lists

If a RADIUS server provides downloadable access lists to Cisco VPN 3000 Series Concentrators as well as to the security appliance, you may need the security appliance to convert wildcard netmask expressions to standard netmask expressions. This is because Cisco VPN 3000 Series Concentrators support wildcard netmask expressions but the security appliance only supports standard netmask expressions. Configuring the security appliance to convert wildcard netmask expressions helps minimize the effects of these differences upon how you configure downloadable access lists on your RADIUS servers. Translation of wildcard netmask expressions means that downloadable access lists written for Cisco VPN 3000 Series Concentrators can be used by the security appliance without altering the configuration of the downloadable access lists on the RADIUS server.

You configure access list netmask conversion on a per server basis, using the **acl-netmask-convert** command, available in the AAA-server configuration mode. For more information about configuring a RADIUS server, see [“Identifying AAA Server Groups and Servers” section on page 13-12](#). For more information about the **acl-netmask-convert** command, see the *Cisco Security Appliance Command Reference*.

## Configuring a RADIUS Server to Download Per-User Access Control List Names

To download a name for an access list that you already created on the security appliance from the RADIUS server when a user authenticates, configure the IETF RADIUS filter-id attribute (attribute number 11) as follows:

```
filter-id=acl_name
```

**Note**

In Cisco Secure ACS, the value for filter-id attributes are specified in boxes in the HTML interface, omitting **filter-id=** and entering only *acl\_name*.

For information about making unique per user the filter-id attribute value, see the documentation for your RADIUS server.

See the [“Adding an Extended Access List” section on page 16-5](#) to create an access list on the security appliance.

## Configuring Accounting for Network Access

The security appliance can send accounting information to a RADIUS or TACACS+ server about any TCP or UDP traffic that passes through the security appliance. If that traffic is also authenticated, then the AAA server can maintain accounting information by username. If the traffic is not authenticated, the AAA server can maintain accounting information by IP address. Accounting information includes when sessions start and stop, username, the number of bytes that pass through the security appliance for the session, the service used, and the duration of each session.

To configure accounting, perform the following steps:

**Step 1** If you want the security appliance to provide accounting data per user, you must enable authentication. For more information, see the [“Enabling Network Access Authentication” section on page 19-3](#). If you want the security appliance to provide accounting data per IP address, enabling authentication is not necessary and you can continue to the next step.

**Step 2** Using the **access-list** command, create an access list that identifies the source addresses and destination addresses of traffic you want accounted. For steps, see the [“Adding an Extended Access List” section on page 16-5](#).

The **permit** ACEs mark matching traffic for authorization, while **deny** entries exclude matching traffic from authorization.

**Note**

If you have configured authentication and want accounting data for all the traffic being authenticated, you can use the same access list you created for use with the **aaa authentication match** command.

**Step 3** To enable accounting, enter the following command:

```
hostname(config)# aaa accounting match acl_name interface_name server_group
```

**Note**

Alternatively, you can use the **aaa accounting include** command (which identifies traffic within the command) but you cannot use both methods in the same configuration. See the *Cisco Security Appliance Command Reference* for more information.

The following commands authenticate, authorize, and account for inside Telnet traffic. Telnet traffic to servers other than 209.165.201.5 can be authenticated alone, but traffic to 209.165.201.5 requires authorization and accounting.

```

hostname(config)# aaa-server AuthOutbound protocol tacacs+
hostname(config-aaa-server-group)# exit
hostname(config)# aaa-server AuthOutbound (inside) host 10.1.1.1
hostname(config-aaa-server-host)# key TACPlusUauthKey
hostname(config-aaa-server-host)# exit
hostname(config)# access-list TELNET_AUTH extended permit tcp any any eq telnet
hostname(config)# access-list SERVER_AUTH extended permit tcp any host 209.165.201.5 eq
telnet
hostname(config)# aaa authentication match TELNET_AUTH inside AuthOutbound
hostname(config)# aaa authorization match SERVER_AUTH inside AuthOutbound
hostname(config)# aaa accounting match SERVER_AUTH inside AuthOutbound

```

## Using MAC Addresses to Exempt Traffic from Authentication and Authorization

The security appliance can exempt from authentication and authorization any traffic from specific MAC addresses. For example, if the security appliance authenticates TCP traffic originating on a particular network but you want to allow unauthenticated TCP connections from a specific server, you would use a MAC exempt rule to exempt from authentication and authorization any traffic from the server specified by the rule.

This feature is particularly useful to exempt devices such as IP phones that cannot respond to authentication prompts.

To use MAC addresses to exempt traffic from authentication and authorization, perform the following steps:

---

**Step 1** To configure a MAC list, enter the following command:

```
hostname(config)# mac-list id {deny | permit} mac macmask
```

Where the *id* argument is the hexadecimal number that you assign to the MAC list. To group a set of MAC addresses, enter the **mac-list** command as many times as needed with the same ID value. Because you can only use one MAC list for AAA exemption, be sure that your MAC list includes all the MAC addresses you want to exempt. You can create multiple MAC lists, but you can only use one at a time.

The order of entries matters, because the packet uses the first entry it matches, as opposed to a best match scenario. If you have a permit entry, and you want to deny an address that is allowed by the permit entry, be sure to enter the deny entry before the permit entry.

The *mac* argument specifies the source MAC address in 12-digit hexadecimal form; that is, nnnn.nnnn.nnnn.

The *macmask* argument specifies the portion of the MAC address that should be used for matching. For example, ffff.ffff.ffff matches the MAC address exactly. ffff.ffff.0000 matches only the first 8 digits.

**Step 2** To exempt traffic for the MAC addresses specified in a particular MAC list, enter the following command:

```
hostname(config)# aaa mac-exempt match id
```

Where *id* is the string identifying the MAC list containing the MAC addresses whose traffic is to be exempt from authentication and authorization. You can only enter one instance of the **aaa mac-exempt** command.

---



The following example bypasses authentication for a single MAC address:

```
hostname(config)# mac-list abc permit 00a0.c95d.0282 ffff.ffff.ffff
hostname(config)# aaa mac-exempt match abc
```

The following entry bypasses authentication for all Cisco IP Phones, which have the hardware ID 0003.E3:

```
hostname(config)# mac-list acd permit 0003.E300.0000 FFFF.FF00.0000
hostname(config)# aaa mac-exempt match acd
```

The following example bypasses authentication for a group of MAC addresses except for 00a0.c95d.02b2. Enter the deny statement before the permit statement, because 00a0.c95d.02b2 matches the permit statement as well, and if it is first, the deny statement will never be matched.

```
hostname(config)# mac-list 1 deny 00a0.c95d.0282 ffff.ffff.ffff
hostname(config)# mac-list 1 permit 00a0.c95d.0000 ffff.ffff.0000
hostname(config)# aaa mac-exempt match 1
```





# CHAPTER 20

## Applying Filtering Services

---

This chapter describes ways to filter web traffic to reduce security risks or prevent inappropriate use. This chapter contains the following sections:

- [Filtering Overview, page 20-1](#)
- [Filtering ActiveX Objects, page 20-2](#)
- [Filtering Java Applets, page 20-3](#)
- [Filtering URLs and FTP Requests with an External Server, page 20-4](#)
- [Viewing Filtering Statistics and Configuration, page 20-9](#)

## Filtering Overview

This section describes how filtering can provide greater control over traffic passing through the security appliance. Filtering can be used in two distinct ways:

- Filtering ActiveX objects or Java applets
- Filtering with an external filtering server

Instead of blocking access altogether, you can remove specific undesirable objects from HTTP traffic, such as ActiveX objects or Java applets, that may pose a security threat in certain situations.

You can also use URL filtering to direct specific traffic to an external filtering server, such as Secure Computing SmartFilter (formerly N2H2) or Websense filtering server. Long URL, HTTPS, and FTP filtering can now be enabled using both Websense and Secure Computing SmartFilter for URL filtering. Filtering servers can block traffic to specific sites or types of sites, as specified by the security policy.



### Note

URL caching will only work if the version of the URL server software from the URL server vendor supports it.

Because URL filtering is CPU-intensive, using an external filtering server ensures that the throughput of other traffic is not affected. However, depending on the speed of your network and the capacity of your URL filtering server, the time required for the initial connection may be noticeably slower when filtering traffic with an external filtering server.

# Filtering ActiveX Objects

This section describes how to apply filtering to remove ActiveX objects from HTTP traffic passing through the firewall. This section includes the following topics:

- [ActiveX Filtering Overview, page 20-2](#)
- [Enabling ActiveX Filtering, page 20-2](#)

## ActiveX Filtering Overview

ActiveX objects may pose security risks because they can contain code intended to attack hosts and servers on a protected network. You can disable ActiveX objects with ActiveX filtering.

ActiveX controls, formerly known as OLE or OCX controls, are components you can insert in a web page or other application. These controls include custom forms, calendars, or any of the extensive third-party forms for gathering or displaying information. As a technology, ActiveX creates many potential problems for network clients including causing workstations to fail, introducing network security problems, or being used to attack servers.

The **filteractivex** command blocks the HTML <object> commands by commenting them out within the HTML web page. ActiveX filtering of HTML files is performed by selectively replacing the <APPLET> and </APPLET> and <OBJECT CLASSID> and </OBJECT> tags with comments. Filtering of nested tags is supported by converting top-level tags to comments.



### Caution

This command also blocks any Java applets, image files, or multimedia objects that are embedded in object tags.

If the <object> or </object> HTML tags split across network packets or if the code in the tags is longer than the number of bytes in the MTU, security appliance cannot block the tag.

ActiveX blocking does not occur when users access an IP address referenced by the **alias** command or for WebVPN traffic.

## Enabling ActiveX Filtering

This section describes how to remove ActiveX objects in HTTP traffic passing through the security appliance. To remove ActiveX objects, enter the following command in global configuration mode:

```
hostname(config)# filteractivex port[-port] local_ip local_mask foreign_ip foreign_mask
```

To use this command, replace *port* with the TCP port to which filtering is applied. Typically, this is port 80, but other values are accepted. The **http** or **url** literal can be used for port 80. You can specify a range of ports by using a hyphen between the starting port number and the ending port number.

The local IP address and mask identify one or more internal hosts that are the source of the traffic to be filtered. The foreign address and mask specify the external destination of the traffic to be filtered.

You can set either address to **0.0.0.0** (or in shortened form, **0**) to specify all hosts. You can use **0.0.0.0** for either mask (or in shortened form, **0**) to specify all hosts.

The following example specifies that ActiveX objects are blocked on all outbound connections:

```
hostname(config)# filteractivex 80 0 0 0 0
```

This command specifies that the ActiveX object blocking applies to web traffic on port 80 from any local host and for connections to any foreign host.

To remove the configuration, use the **no** form of the command, as in the following example:

```
hostname(config)# no filteractivex 80 0 0 0 0
```

## Filtering Java Applets

This section describes how to apply filtering to remove Java applets from HTTP traffic passing through the firewall. Java applets may pose security risks because they can contain code intended to attack hosts and servers on a protected network. You can remove Java applets with the **filter java** command.

The **filter java** command filters out Java applets that return to the security appliance from an outbound connection. The user still receives the HTML page, but the web page source for the applet is commented out so that the applet cannot execute. The **filter java** command does not filter WebVPN traffic.



### Note

Use the **filteractivex** command to remove Java applets that are embedded in <object> tags.

To remove Java applets in HTTP traffic passing through the firewall, enter the following command in global configuration mode:

```
hostname(config)# filter java port[-port] local_ip local_mask foreign_ip foreign_mask
```

To use this command, replace *port* with the TCP port to which filtering is applied. Typically, this is port 80, but other values are accepted. The **http** or **url** literal can be used for port 80. You can specify a range of ports by using a hyphen between the starting port number and the ending port number.

The local IP address and mask identify one or more internal hosts that are the source of the traffic to be filtered. The foreign address and mask specify the external destination of the traffic to be filtered.

You can set either address to **0.0.0.0** (or in shortened form, **0**) to specify all hosts. You can use **0.0.0.0** for either mask (or in shortened form, **0**) to specify all hosts.

You can set either address to **0.0.0.0** (or in shortened form, **0**) to specify all hosts. You can use **0.0.0.0** for either mask (or in shortened form, **0**) to specify all hosts.

The following example specifies that Java applets are blocked on all outbound connections:

```
hostname(config)# filter java 80 0 0 0 0
```

This command specifies that the Java applet blocking applies to web traffic on port 80 from any local host and for connections to any foreign host.

The following example blocks downloading of Java applets to a host on a protected network:

```
hostname(config)# filter java http 192.168.3.3 255.255.255.255 0 0
```

This command prevents host 192.168.3.3 from downloading Java applets.

To remove the configuration, use the **no** form of the command, as in the following example:

```
hostname(config)# no filter java http 192.168.3.3 255.255.255.255 0 0
```

# Filtering URLs and FTP Requests with an External Server

This section describes how to filter URLs and FTP requests with an external server. This section includes the following topics:

- [URL Filtering Overview, page 20-4](#)
- [Identifying the Filtering Server, page 20-4](#)
- [Buffering the Content Server Response, page 20-6](#)
- [Caching Server Addresses, page 20-6](#)
- [Filtering HTTP URLs, page 20-7](#)
- [Filtering HTTPS URLs, page 20-8](#)
- [Filtering FTP Requests, page 20-9](#)

## URL Filtering Overview

You can apply filtering to connection requests originating from a more secure network to a less secure network. Although you can use ACLs to prevent outbound access to specific content servers, managing usage this way is difficult because of the size and dynamic nature of the Internet. You can simplify configuration and improve security appliance performance by using a separate server running one of the following Internet filtering products:

- Websense Enterprise for filtering HTTP, HTTPS, and FTP.
- Secure Computing SmartFilter (formerly N2H2) for filtering HTTP, HTTPS, FTP, and long URL filtering.

**Note**

---

URL caching will only work if the version of the URL server software from the URL server vendor supports it.

---

Although security appliance performance is less affected when using an external server, users may notice longer access times to websites or FTP servers when the filtering server is remote from the security appliance.

When filtering is enabled and a request for content is directed through the security appliance, the request is sent to the content server and to the filtering server at the same time. If the filtering server allows the connection, the security appliance forwards the response from the content server to the originating client. If the filtering server denies the connection, the security appliance drops the response and sends a message or return code indicating that the connection was not successful.

If user authentication is enabled on the security appliance, then the security appliance also sends the user name to the filtering server. The filtering server can use user-specific filtering settings or provide enhanced reporting regarding usage.

## Identifying the Filtering Server

You can identify up to four filtering servers per context. The security appliance uses the servers in order until a server responds. You can only configure a single type of server (Websense or Secure Computing SmartFilter ) in your configuration.

**Note**

You must add the filtering server before you can configure filtering for HTTP or HTTPS with the **filter** command. If you remove the filtering servers from the configuration, then all **filter** commands are also removed.

Identify the address of the filtering server using the **url-server** command:

For Websense:

```
hostname(config)# url-server (if_name) host local_ip [timeout seconds] [protocol TCP | UDP
version [1|4] [connections num_conns]]
```

For Secure Computing SmartFilter (formerly N2H2):

```
hostname(config)# url-server (if_name) vendor {secure-computing | n2h2} host
<local_ip> [port <number>] [timeout <seconds>] [protocol {TCP [connections <number>]} |
UDP]
```

where *<if\_name>* is the name of the security appliance interface connected to the filtering server (the default is inside).

For the **vendor** {secure-computing | n2h2}, you can use 'secure-computing' as a vendor string, however, 'n2h2' is acceptable for backward compatibility. When the configuration entries are generated, 'secure-computing' is saved as the vendor string.

The **host** *<local\_ip>* is the IP address of the URL filtering server.

The **port** *<number>* is the Secure Computing SmartFilter server port number of the filtering server; the security appliance also listens for UDP replies on this port.

**Note**

The default port is 4005. This is the default port used by the Secure Computing SmartFilter server to communicate to the security appliance via TCP or UDP. For information on changing the default port, please refer to the *Filtering by N2H2 Administrator's Guide*.

The **timeout** *<seconds>* is the number of seconds the security appliance should keep trying to connect to the filtering server.

The **connections** *<number>* is the number of tries to attempt to make a connection between the host and server.

For example, to identify a single Websense filtering server, enter the following command:

```
hostname(config)# url-server (perimeter) host 10.0.1.1 protocol TCP version 4
```

This identifies a Websense filtering server with the IP address 10.0.1.1 on a perimeter interface of the security appliance. Version 4, which is enabled in this example, is recommended by Websense because it supports caching.

To identify redundant Secure Computing SmartFilter servers, enter the following commands:

```
hostname(config)# url-server (perimeter) vendor n2h2 host 10.0.1.1
hostname(config)# url-server (perimeter) vendor n2h2 host 10.0.1.2
```

This identifies two Sentian filtering servers, both on a perimeter interface of the security appliance.

## Buffering the Content Server Response

When a user issues a request to connect to a content server, the security appliance sends the request to the content server and to the filtering server at the same time. If the filtering server does not respond before the content server, the server response is dropped. This delays the web server response from the point of view of the web client because the client must reissue the request.

By enabling the HTTP response buffer, replies from web content servers are buffered and the responses are forwarded to the requesting client if the filtering server allows the connection. This prevents the delay that might otherwise occur.

To configure buffering for responses to HTTP or FTP requests, perform the following steps:

- Step 1** To enable buffering of responses for HTTP or FTP requests that are pending a response from the filtering server, enter the following command:

```
hostname(config)# url-block block block-buffer-limit
```

Replace *block-buffer* with the maximum number of HTTP responses that can be buffered while awaiting responses from the url-server.



**Note** Buffering URLs longer than 3072 bytes are not supported.

- Step 2** To configure the maximum memory available for buffering pending URLs (and for buffering long URLs), enter the following command:

```
hostname(config)# url-block mempool-size memory-pool-size
```

Replace *memory-pool-size* with a value from 2 to 10240 for a maximum memory allocation of 2 KB to 10 MB.

## Caching Server Addresses

After a user accesses a site, the filtering server can allow the security appliance to cache the server address for a certain amount of time, as long as every site hosted at the address is in a category that is permitted at all times. Then, when the user accesses the server again, or if another user accesses the server, the security appliance does not need to consult the filtering server again.



**Note**

Requests for cached IP addresses are not passed to the filtering server and are not logged. As a result, this activity does not appear in any reports. You can accumulate Websense run logs before using the **url-cache** command.

Use the **url-cache** command if needed to improve throughput, as follows:

```
hostname(config)# url-cache dst | src_dst size
```

Replace *size* with a value for the cache size within the range 1 to 128 (KB).

Use the **dst** keyword to cache entries based on the URL destination address. Select this mode if all users share the same URL filtering policy on the Websense server.



Use the **src\_dst** keyword to cache entries based on both the source address initiating the URL request as well as the URL destination address. Select this mode if users do not share the same URL filtering policy on the Websense server.

## Filtering HTTP URLs

This section describes how to configure HTTP filtering with an external filtering server. This section includes the following topics:

- [Configuring HTTP Filtering, page 20-7](#)
- [Enabling Filtering of Long HTTP URLs, page 20-7](#)
- [Truncating Long HTTP URLs, page 20-7](#)
- [Exempting Traffic from Filtering, page 20-8](#)

## Configuring HTTP Filtering

You must identify and enable the URL filtering server before enabling HTTP filtering.

When the filtering server approves an HTTP connection request, the security appliance allows the reply from the web server to reach the originating client. If the filtering server denies the request, the security appliance redirects the user to a block page, indicating that access was denied.

To enable HTTP filtering, enter the following command:

```
hostname(config)# filter url [http | port[-port] local_ip local_mask foreign_ip
foreign_mask] [allow] [proxy-block]
```

Replace *port* with one or more port numbers if a different port than the default port for HTTP (80) is used. Replace *local\_ip* and *local\_mask* with the IP address and subnet mask of a user or subnetwork making requests. Replace *foreign\_ip* and *foreign\_mask* with the IP address and subnet mask of a server or subnetwork responding to requests.

The **allow** option causes the security appliance to forward HTTP traffic without filtering when the primary filtering server is unavailable. Use the **proxy-block** command to drop all requests to proxy servers.

## Enabling Filtering of Long HTTP URLs

By default, the security appliance considers an HTTP URL to be a long URL if it is greater than 1159 characters. You can increase the maximum length allowed.

Configure the maximum size of a single URL with the following command:

```
hostname(config)# url-block url-size long-url-size
```

Replace *long-url-size* with the maximum size in KB for each long URL being buffered. For Websense, this is a value from 2 to 4 for a maximum URL size of 2 KB to 4 KB; for Secure Computing, this is a value between 2 to 3 for a maximum URL size of 2 KB to 3 KB. The default value is 2.

## Truncating Long HTTP URLs

By default, if a URL exceeds the maximum permitted size, then it is dropped. To avoid this, you can set the security appliance to truncate a long URL by entering the following command:

```
hostname(config)# filter url [longurl-truncate | longurl-deny | cgi-truncate]
```

The **longurl-truncate** option causes the security appliance to send only the hostname or IP address portion of the URL for evaluation to the filtering server when the URL is longer than the maximum length permitted. Use the **longurl-deny** option to deny outbound URL traffic if the URL is longer than the maximum permitted.

Use the **cgi-truncate** option to truncate CGI URLs to include only the CGI script location and the script name without any parameters. Many long HTTP requests are CGI requests. If the parameters list is very long, waiting and sending the complete CGI request including the parameter list can use up memory resources and affect firewall performance.

## Exempting Traffic from Filtering

To exempt specific traffic from filtering, enter the following command:

```
hostname(config)# filter url except source_ip source_mask dest_ip dest_mask
```

For example, the following commands cause all HTTP requests to be forwarded to the filtering server except for those from 10.0.2.54.

```
hostname(config)# filter url http 0 0 0 0
hostname(config)# filter url except 10.0.2.54 255.255.255.255 0 0
```

## Filtering HTTPS URLs

You must identify and enable the URL filtering server before enabling HTTPS filtering.



### Note

Websense and Smartfilter currently support HTTPS; older versions of Secure Computing SmartFilter (formerly N2H2) did not support HTTPS filtering.

Because HTTPS content is encrypted, the security appliance sends the URL lookup without directory and filename information. When the filtering server approves an HTTPS connection request, the security appliance allows the completion of SSL connection negotiation and allows the reply from the web server to reach the originating client. If the filtering server denies the request, the security appliance prevents the completion of SSL connection negotiation. The browser displays an error message such as “The Page or the content cannot be displayed.”



### Note

The security appliance does not provide an authentication prompt for HTTPS, so a user must authenticate with the security appliance using HTTP or FTP before accessing HTTPS servers.

To enable HTTPS filtering, enter the following command:

```
hostname(config)# filter https port[-port] localIP local_mask foreign_IP foreign_mask
[allow]
```

Replace *port[-port]* with a range of port numbers if a different port than the default port for HTTPS (443) is used.

Replace *local\_ip* and *local\_mask* with the IP address and subnet mask of a user or subnetwork making requests.

Replace *foreign\_ip* and *foreign\_mask* with the IP address and subnet mask of a server or subnetwork responding to requests.

The **allow** option causes the security appliance to forward HTTPS traffic without filtering when the primary filtering server is unavailable.

## Filtering FTP Requests

You must identify and enable the URL filtering server before enabling FTP filtering.



### Note

Websense and Smartfilter currently support FTP; older versions of Secure Computing SmartFilter (formerly known as N2H2) did not support FTP filtering.

When the filtering server approves an FTP connection request, the security appliance allows the successful FTP return code to reach originating client. For example, a successful return code is “250: CWD command successful.” If the filtering server denies the request, alters the FTP return code to show that the connection was denied. For example, the security appliance changes code 250 to “550 Requested file is prohibited by URL filtering policy.”

To enable FTP filtering, enter the following command:

```
hostname(config)# filter ftp port[-port] localIP local_mask foreign_IP foreign_mask
[allow] [interact-block]
```

Replace *port[-port]* with a range of port numbers if a different port than the default port for FTP (21) is used.

Replace *local\_ip* and *local\_mask* with the IP address and subnet mask of a user or subnetwork making requests.

Replace *foreign\_ip* and *foreign\_mask* with the IP address and subnet mask of a server or subnetwork responding to requests.

The **allow** option causes the security appliance to forward HTTPS traffic without filtering when the primary filtering server is unavailable.

Use the **interact-block** option to prevent interactive FTP sessions that do not provide the entire directory path. An interactive FTP client allows the user to change directories without typing the entire path. For example, the user might enter **cd ./files** instead of **cd /public/files**.

## Viewing Filtering Statistics and Configuration

This section describes how to monitor filtering statistics. This section includes the following topics:

- [Viewing Filtering Server Statistics, page 20-10](#)
- [Viewing Buffer Configuration and Statistics, page 20-11](#)
- [Viewing Caching Statistics, page 20-11](#)
- [Viewing Filtering Performance Statistics, page 20-11](#)
- [Viewing Filtering Configuration, page 20-12](#)

## Viewing Filtering Server Statistics

To show information about the filtering server, enter the following command:

```
hostname# show running-config url-server
```

The following is sample output from the **show running-config url-server** command:

```
hostname# show running-config url-server
url-server (outside) vendor n2h2 host 128.107.254.202 port 4005 timeout 5 protocol TCP
```

To show information about the filtering server or to show statistics, enter the following command:

The following is sample output from the **show running-config url-server statistics** command, which shows filtering statistics:

```
hostname# show running-config url-server statistics
```

Global Statistics:

```

URLs total/allowed/denied 13/3/10
URLs allowed by cache/server 0/3
URLs denied by cache/server 0/10
HTTPSs total/allowed/denied 138/137/1
HTTPSs allowed by cache/server 0/137
HTTPSs denied by cache/server 0/1
FTPs total/allowed/denied 0/0/0
FTPs allowed by cache/server 0/0
FTPs denied by cache/server 0/0
Requests dropped 0
Server timeouts/retries 0/0
Processed rate average 60s/300s 0/0 requests/second
Denied rate average 60s/300s 0/0 requests/second
Dropped rate average 60s/300s 0/0 requests/second
```

Server Statistics:

```

10.125.76.20 UP
 Vendor websense
 Port 15868
 Requests total/allowed/denied 151/140/11
 Server timeouts/retries 0/0
 Responses received 151
 Response time average 60s/300s 0/0
```

URL Packets Sent and Received Stats:

```

Message Sent Received
STATUS_REQUEST 1609 1601
LOOKUP_REQUEST 1526 1526
LOG_REQUEST 0 NA
```

Errors:

```

RFC noncompliant GET method 0
URL buffer update failure 0
```

## Viewing Buffer Configuration and Statistics

The **show running-config url-block** command displays the number of packets held in the url-block buffer and the number (if any) dropped due to exceeding the buffer limit or retransmission.

The following is sample output from the **show running-config url-block** command:

```
hostname# show running-config url-block
url-block url-mempool 128
url-block url-size 4
url-block block 128
```

This shows the configuration of the URL block buffer.

The following is sample output from the **show url-block block statistics** command:

```
hostname# show running-config url-block block statistics

URL Pending Packet Buffer Stats with max block 128

Cumulative number of packets held: 896
Maximum number of packets held (per URL): 3
Current number of packets held (global): 38
Packets dropped due to
 exceeding url-block buffer limit: 7546
 HTTP server retransmission: 10
Number of packets released back to client: 0
```

This shows the URL block statistics.

## Viewing Caching Statistics

The following is sample output from the **show url-cache stats** command:

```
hostname# show url-cache stats
URL Filter Cache Stats

Size : 128KB
Entries : 1724
In Use : 456
Lookups : 45
Hits : 8
```

This shows how the cache is used.

## Viewing Filtering Performance Statistics

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

```
hostname# show perfmon
PERFMON STATS: Current Average
Xlates 0/s 0/s
Connections 0/s 2/s
TCP Conns 0/s 2/s
UDP Conns 0/s 0/s
URL Access 0/s 2/s
URL Server Req 0/s 3/s
TCP Fixup 0/s 0/s
TCPIntercept 0/s 0/s
HTTP Fixup 0/s 3/s
```

|             |     |     |
|-------------|-----|-----|
| FTP Fixup   | 0/s | 0/s |
| AAA Authen  | 0/s | 0/s |
| AAA Author  | 0/s | 0/s |
| AAA Account | 0/s | 0/s |

This shows URL filtering performance statistics, along with other performance statistics. The filtering statistics are shown in the URL Access and URL Server Req rows.

## Viewing Filtering Configuration

The following is sample output from the **show running-config filter** command:

```
hostname# show running-config filter
filter url http 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
```



# CHAPTER 21

## Using Modular Policy Framework

---

This chapter describes how to use Modular Policy Framework to create security policies for TCP and general connection settings, inspections, IPS, CSC, and QoS.

This chapter includes the following sections:

- [Modular Policy Framework Overview, page 21-1](#)
- [Identifying Traffic \(Layer 3/4 Class Map\), page 21-4](#)
- [Configuring Special Actions for Application Inspections \(Inspection Policy Map\), page 21-7](#)
- [Defining Actions \(Layer 3/4 Policy Map\), page 21-15](#)
- [Applying Actions to an Interface \(Service Policy\), page 21-21](#)
- [Modular Policy Framework Examples, page 21-21](#)

## Modular Policy Framework Overview

Modular Policy Framework provides a consistent and flexible way to configure security appliance features. For example, you can use Modular Policy Framework to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. This section includes the following topics:

- [Modular Policy Framework Features, page 21-1](#)
- [Modular Policy Framework Configuration Overview, page 21-2](#)
- [Default Global Policy, page 21-3](#)

## Modular Policy Framework Features

Modular Policy Framework supports the following features:

- QoS input policing
- TCP normalization, TCP and UDP connection limits and timeouts, and TCP sequence number randomization
- CSC
- Application inspection
- IPS
- QoS output policing


- QoS standard priority queue
- QoS traffic shaping, hierarchical priority queue

## Modular Policy Framework Configuration Overview


Configuring Modular Policy Framework consists of the following tasks:

1. Identify the traffic on which you want to perform Modular Policy Framework actions by creating Layer 3/4 class maps. For example, you might want to perform actions on all traffic that passes through the security appliance; or you might only want to perform certain actions on traffic from 10.1.1.0/24 to any destination address.

Layer 3/4 Class Map



Layer 3/4 Class Map




241506

See the [“Identifying Traffic \(Layer 3/4 Class Map\)”](#) section on page 21-4.

2. If one of the actions you want to perform is application inspection, and you want to perform additional actions on some inspection traffic, then create an inspection policy map. The inspection policy map identifies the traffic and specifies what to do with it. For example, you might want to drop all HTTP requests with a body length greater than 1000 bytes.

Inspection Policy Map Actions

Inspection Class Map/  
Match Commands




241507

You can create a self-contained inspection policy map that identifies the traffic directly with **match** commands, or you can create an inspection class map for reuse or for more complicated matching. See the [“Defining Actions in an Inspection Policy Map”](#) section on page 21-8 and the [“Identifying Traffic in an Inspection Class Map”](#) section on page 21-11.


3. If you want to match text with a regular expression within inspected packets, you can create a regular expression or a group of regular expressions (a regular expression class map). Then, when you define the traffic to match for the inspection policy map, you can call on an existing regular expression. For example, you might want to drop all HTTP requests with a URL including the text “example.com.”





See the “Creating a Regular Expression” section on page 21-12 and the “Creating a Regular Expression Class Map” section on page 21-14.

4. Define the actions you want to perform on each Layer 3/4 class map by creating a Layer 3/4 policy map. Then, determine on which interfaces you want to apply the policy map using a service policy.



See the “Defining Actions (Layer 3/4 Policy Map)” section on page 21-15 and the “Applying Actions to an Interface (Service Policy)” section on page 21-21.

## Default Global Policy

By default, the configuration includes a policy that matches all default application inspection traffic and applies certain inspections to the traffic on all interfaces (a global policy). Not all inspections are enabled by default. You can only apply one global policy, so if you want to alter the global policy, you need to either edit the default policy or disable it and apply a new one. (An interface policy overrides the global policy for a particular feature.)

The default policy configuration includes the following commands:

```
class-map inspection_default
 match default-inspection-traffic
policy-map type inspect dns preset_dns_map
 parameters
 message-length maximum 512
policy-map global_policy
 class inspection_default
 inspect dns preset_dns_map
 inspect ftp
 inspect h323 h225
 inspect h323 ras
 inspect rsh
 inspect rtsp
 inspect esmtp
 inspect sqlnet
 inspect skinny
 inspect sunrpc
 inspect xdmcp
 inspect sip
 inspect netbios
 inspect tftp
service-policy global_policy global
```

## Identifying Traffic (Layer 3/4 Class Map)

A Layer 3/4 class map identifies Layer 3 and 4 traffic to which you want to apply actions. The maximum number of Layer 3/4 class maps is 255 in single mode or per context in multiple mode. You can create multiple Layer 3/4 class maps for each Layer 3/4 policy map. You can create the following types of class maps:

- [Default Class Maps, page 21-4](#)
- [Creating a Layer 3/4 Class Map for Through Traffic, page 21-5](#)
- [Creating a Layer 3/4 Class Map for Management Traffic, page 21-7](#)

## Default Class Maps

The configuration includes a default Layer 3/4 class map that the security appliance uses in the default global policy. It is called **inspection\_default** and matches the default inspection traffic:

```
class-map inspection_default
 match default-inspection-traffic
```

Another class map that exists in the default configuration is called **class-default**, and it matches all traffic:

```
class-map class-default
 match any
```

This class map appears at the end of all Layer 3/4 policy maps and essentially tells the security appliance to not perform any actions on all other traffic. You can use the **class-default** class map if desired, rather than making your own **match any** class map. In fact, some features are only available for **class-default**, such as QoS traffic shaping.

## Creating a Layer 3/4 Class Map for Through Traffic

A Layer 3/4 class map matches traffic based on protocols, ports, IP addresses and other Layer 3 or 4 attributes.

To define a Layer 3/4 class map, perform the following steps:

**Step 1** Create a Layer 3/4 class map by entering the following command:

```
hostname(config)# class-map class_map_name
hostname(config-cmap)#
```

Where *class\_map\_name* is a string up to 40 characters in length. The name “class-default” is reserved. All types of class maps use the same name space, so you cannot reuse a name already used by another type of class map. The CLI enters class-map configuration mode.

**Step 2** (Optional) Add a description to the class map by entering the following command:

```
hostname(config-cmap)# description string
```

**Step 3** Define the traffic to include in the class by matching one of the following characteristics. Unless otherwise specified, you can include only one **match** command in the class map.

- Any traffic—The class map matches all traffic.

```
hostname(config-cmap)# match any
```

- Access list—The class map matches traffic specified by an extended access list. If the security appliance is operating in transparent firewall mode, you can use an EtherType access list.

```
hostname(config-cmap)# match access-list access_list_name
```

For more information about creating access lists, see the [“Adding an Extended Access List” section on page 16-5](#) or the [“Adding an EtherType Access List” section on page 16-8](#).

For information about creating access lists with NAT, see the [“IP Addresses Used for Access Lists When You Use NAT” section on page 16-3](#).

- TCP or UDP destination ports—The class map matches a single port or a contiguous range of ports.

```
hostname(config-cmap)# match port {tcp | udp} {eq port_num | range port_num port_num}
```



### Tip

For applications that use multiple, non-contiguous ports, use the **match access-list** command and define an ACE to match each port.

For a list of ports you can specify, see the [“TCP and UDP Ports” section on page D-11](#).

For example, enter the following command to match TCP packets on port 80 (HTTP):

```
hostname(config-cmap)# match tcp eq 80
```

- Default traffic for inspection—The class map matches the default TCP and UDP ports used by all applications that the security appliance can inspect.

```
hostname(config-cmap)# match default-inspection-traffic
```

See the [“Default Inspection Policy” section on page 25-3](#) for a list of default ports. The security appliance includes a default global policy that matches the default inspection traffic, and applies common inspections to the traffic on all interfaces. Not all applications whose ports are included in the **match default-inspection-traffic** command are enabled by default in the policy map.

You can specify a **match access-list** command along with the **match default-inspection-traffic** command to narrow the matched traffic. Because the **match default-inspection-traffic** command specifies the ports to match, any ports in the access list are ignored.

- DSCP value in an IP header—The class map matches up to eight DSCP values.

```
hostname(config-cmap)# match dscp value1 [value2] [...] [value8]
```

For example, enter the following:

```
hostname(config-cmap)# match dscp af43 cs1 ef
```

- Precedence—The class map matches up to four precedence values, represented by the Type of Service (TOS) byte in the IP header.

```
hostname(config-cmap)# match precedence value1 [value2] [value3] [value4]
```

where *value1* through *value4* can be 0 to 7, corresponding to the possible precedences.

- RTP traffic—The class map matches RTP traffic.

```
hostname(config-cmap)# match rtp starting_port range
```

The *starting\_port* specifies an even-numbered UDP destination port between 2000 and 65534. The *range* specifies the number of additional UDP ports to match above the *starting\_port*, between 0 and 16383.

- Tunnel group traffic—The class map matches traffic for a tunnel group to which you want to apply QoS.

```
hostname(config-cmap)# match tunnel-group name
```

You can also specify one other **match** command to refine the traffic match. You can specify any of the preceding commands, except for the **match any**, **match access-list**, or **match default-inspection-traffic** commands. Or you can enter the following command to police each flow:

```
hostname(config-cmap)# match flow ip destination address
```

All traffic going to a unique IP destination address is considered a flow.

The following is an example for the **class-map** command:

```
hostname(config)# access-list udp permit udp any any
hostname(config)# access-list tcp permit tcp any any
hostname(config)# access-list host_foo permit ip any 10.1.1.1 255.255.255.255

hostname(config)# class-map all_udp
hostname(config-cmap)# description "This class-map matches all UDP traffic"
hostname(config-cmap)# match access-list udp

hostname(config-cmap)# class-map all_tcp
hostname(config-cmap)# description "This class-map matches all TCP traffic"
hostname(config-cmap)# match access-list tcp

hostname(config-cmap)# class-map all_http
hostname(config-cmap)# description "This class-map matches all HTTP traffic"
hostname(config-cmap)# match port tcp eq http

hostname(config-cmap)# class-map to_server
hostname(config-cmap)# description "This class-map matches all traffic to server 10.1.1.1"
hostname(config-cmap)# match access-list host_foo
```

## Creating a Layer 3/4 Class Map for Management Traffic

For management traffic to the security appliance, you might want to perform actions specific to this kind of traffic. You can specify a management class map that can match TCP or UDP ports. The types of actions available for a management class map in the policy map are specialized for management traffic. Namely, this type of class map lets you inspect RADIUS accounting traffic.

To create a class map for management traffic to the security appliance, perform the following steps:

---

**Step 1** Create a class map by entering the following command:

```
hostname(config)# class-map type management class_map_name
hostname(config-cmap)#
```

Where *class\_map\_name* is a string up to 40 characters in length. The name “class-default” is reserved. All types of class maps use the same name space, so you cannot reuse a name already used by another type of class map. The CLI enters class-map configuration mode.

**Step 2** (Optional) Add a description to the class map by entering the following command:

```
hostname(config-cmap)# description string
```

**Step 3** Define the traffic to include in the class by matching the TCP or UDP port. You can include only one **match** command in the class map.

```
hostname(config-cmap)# match port {tcp | udp} {eq port_num | range port_num port_num}
```

For a list of ports you can specify, see the [“TCP and UDP Ports” section on page D-11](#).

For example, enter the following command to match TCP packets on port 10000:

```
hostname(config-cmap)# match tcp eq 10000
```

---

## Configuring Special Actions for Application Inspections (Inspection Policy Map)

Modular Policy Framework lets you configure special actions for many application inspections. When you enable an inspection engine in the Layer 3/4 policy map, you can also optionally enable actions as defined in an *inspection policy map*. When the inspection policy map matches traffic within the Layer 3/4 class map for which you have defined an inspection action, then that subset of traffic will be acted upon as specified (for example, dropped or rate-limited).

This section includes the following topics:

- [Inspection Policy Map Overview, page 21-8](#)
- [Defining Actions in an Inspection Policy Map, page 21-8](#)
- [Identifying Traffic in an Inspection Class Map, page 21-11](#)
- [Creating a Regular Expression, page 21-12](#)
- [Creating a Regular Expression Class Map, page 21-14](#)

## Inspection Policy Map Overview

See the “[Configuring Application Inspection](#)” section on page 25-5 for a list of applications that support inspection policy maps.

An inspection policy map consists of one or more of the following elements. The exact options available for an inspection policy map depends on the application.

- Traffic matching command—You can define a traffic matching command directly in the inspection policy map to match application traffic to criteria specific to the application, such as a URL string, for which you then enable actions.
  - Some traffic matching commands can specify regular expressions to match text inside a packet. Be sure to create and test the regular expressions before you configure the policy map, either singly or grouped together in a regular expression class map.
- Inspection class map—(Not available for all applications. See the CLI help for a list of supported applications.) An inspection class map includes traffic matching commands that match application traffic with criteria specific to the application, such as a URL string. You then identify the class map in the policy map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspection policy map is that you can create more complex match criteria and you can reuse class maps.
  - Some traffic matching commands can specify regular expressions to match text inside a packet. Be sure to create and test the regular expressions before you configure the policy map, either singly or grouped together in a regular expression class map.
- Parameters—Parameters affect the behavior of the inspection engine.

The default inspection policy map configuration includes the following commands, which sets the maximum message length for DNS packets to be 512 bytes:

```
policy-map type inspect dns preset_dns_map
 parameters
 message-length maximum 512
```



### Note

There are other default inspection policy maps such as **policy-map type inspect esmtp \_default\_esmtp\_map**. These default policy maps are created implicitly by the command **inspect protocol**. For example, **inspect esmtp** implicitly uses the policy map “\_default\_esmtp\_map.” All the default policy maps can be shown by using the **show running-config all policy-map** command.

## Defining Actions in an Inspection Policy Map

When you enable an inspection engine in the Layer 3/4 policy map, you can also optionally enable actions as defined in an inspection policy map.

To create an inspection policy map, perform the following steps:

- Step 1** To create the HTTP inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect application policy_map_name
hostname(config-pmap)#
```

See the “[Configuring Application Inspection](#)” section on page 25-5 for a list of applications that support inspection policy maps.

The *policy\_map\_name* argument is the name of the policy map up to 40 characters in length. All types of policy maps use the same name space, so you cannot reuse a name already used by another type of policy map. The CLI enters policy-map configuration mode.

**Step 2** To apply actions to matching traffic, perform the following steps:

a. Specify the traffic on which you want to perform actions using one of the following methods:

- Specify the inspection class map that you created in the “[Identifying Traffic in an Inspection Class Map](#)” section on page 21-11 by entering the following command:

```
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

- Specify traffic directly in the policy map using one of the **match** commands described for each application in [Chapter 25, “Configuring Application Layer Protocol Inspection.”](#) If you use a **match not** command, then any traffic that matches the criterion in the **match not** command does not have the action applied.

b. Specify the action you want to perform on the matching traffic by entering the following command:

```
hostname(config-pmap-c)# {[drop [send-protocol-error] |
drop-connection [send-protocol-error] | mask | reset] [log] | rate-limit message_rate}
```

Not all options are available for each application. Other actions specific to the application might also be available. See [Chapter 25, “Configuring Application Layer Protocol Inspection,”](#) for the exact options available.

The **drop** keyword drops all packets that match.

The **send-protocol-error** keyword sends a protocol error message.

The **drop-connection** keyword drops the packet and closes the connection.

The **mask** keyword masks out the matching portion of the packet.

The **reset** keyword drops the packet, closes the connection, and sends a TCP reset to the server and/or client.

The **log** keyword, which you can use alone or with one of the other keywords, sends a system log message.

The **rate-limit** *message\_rate* argument limits the rate of messages.



#### Note

You can specify multiple **class** or **match** commands in the policy map.

If a packet matches multiple different **match** or **class** commands, then the order in which the security appliance applies the actions is determined by internal security appliance rules, and not by the order they are added to the policy map. The internal rules are determined by the application type and the logical progression of parsing a packet, and are not user-configurable. For example for HTTP traffic, parsing a Request Method field precedes parsing the Header Host Length field; an action for the Request Method field occurs before the action for the Header Host Length field. For example, the following match commands can be entered in any order, but the **match request method get** command is matched first.

```
match request header host length gt 100
 reset
match request method get
 log
```

If an action drops a packet, then no further actions are performed in the inspection policy map. For example, if the first action is to reset the connection, then it will never match any further **match** or **class** commands. If the first action is to log the packet, then a second action, such as resetting the connection, can occur. (You can configure both the **reset** (or **drop-connection**, and so on.) and the **log** action for the same **match** or **class** command, in which case the packet is logged before it is reset for a given match.)

If a packet matches multiple **match** or **class** commands that are the same, then they are matched in the order they appear in the policy map. For example, for a packet with the header length of 1001, it will match the first command below, and be logged, and then will match the second command and be reset. If you reverse the order of the two **match** commands, then the packet will be dropped and the connection reset before it can match the second **match** command; it will never be logged.

```
match request header length gt 100
 log
match request header length gt 1000
 reset
```

A class map is determined to be the same type as another class map or **match** command based on the lowest priority **match** command in the class map (the priority is based on the internal rules). If a class map has the same type of lowest priority **match** command as another class map, then the class maps are matched according to the order they are added to the policy map. If the lowest priority command for each class map is different, then the class map with the higher priority **match** command is matched first. For example, the following three class maps contain two types of **match** commands: **match request-cmd** (higher priority) and **match filename** (lower priority). The ftp3 class map includes both commands, but it is ranked according to the lowest priority command, **match filename**. The ftp1 class map includes the highest priority command, so it is matched first, regardless of the order in the policy map. The ftp3 class map is ranked as being of the same priority as the ftp2 class map, which also contains the **match filename** command. They are matched according to the order in the policy map: ftp3 and then ftp2.

```
class-map inspect type ftp ftp1
 match request-cmd get
class-map inspect type ftp ftp2
 match filename regex abc
class-map inspect type ftp ftp3
 match request-cmd get
 match filename regex abc

policy-map type inspect ftp ftp
 class ftp3
 log
 class ftp2
 log
 class ftp1
 log
```

---

**Step 3** To configure parameters that affect the inspection engine, enter the following command:

```
hostname(config-pmap)# parameters
hostname(config-pmap-p)#
```

The CLI enters parameters configuration mode. For the parameters available for each application, see [Chapter 25, “Configuring Application Layer Protocol Inspection.”](#)

---

The following is an example of an HTTP inspection policy map and the related class maps. This policy map is activated by the Layer 3/4 policy map, which is enabled by the service policy.

```
hostname(config)# regex url_example example.com
```



```

hostname(config)# regex url_example2 example2.com
hostname(config)# class-map type regex match-any URLs
hostname(config-cmap)# match regex url_example
hostname(config-cmap)# match regex url_example2

hostname(config-cmap)# class-map type inspect http match-all http-traffic
hostname(config-cmap)# match req-resp content-type mismatch
hostname(config-cmap)# match request body length gt 1000
hostname(config-cmap)# match not request uri regex class URLs

hostname(config-cmap)# policy-map type inspect http http-map1
hostname(config-pmap)# class http-traffic
hostname(config-pmap-c)# drop-connection log
hostname(config-pmap-c)# match req-resp content-type mismatch
hostname(config-pmap-c)# reset log
hostname(config-pmap-c)# parameters
hostname(config-pmap-p)# protocol-violation action log

hostname(config-pmap-p)# policy-map test
hostname(config-pmap)# class test (a Layer 3/4 class map not shown)
hostname(config-pmap-c)# inspect http http-map1

hostname(config-pmap-c)# service-policy test interface outside

```

## Identifying Traffic in an Inspection Class Map

This type of class map allows you to match criteria that is specific to an application. For example, for DNS traffic, you can match the domain name in a DNS query.



### Note

Not all applications support inspection class maps. See the CLI help for a list of supported applications.

A class map groups multiple traffic matches. Traffic must match *all* of the match criteria to match the class map. You can alternatively identify the traffic you want to match directly in the policy map. The difference between creating a class map and defining the traffic match directly in the inspection policy map is that the class map lets you group multiple matches, and you can reuse class maps. For the traffic that you identify in this class map, you can specify actions such as dropping, resetting, and/or logging the connection in the inspection policy map. If you want to perform different actions on different types of traffic, you should identify the traffic directly in the policy map.

To define an inspection class map, perform the following steps:

**Step 1** Create a class map by entering the following command:

```

hostname(config)# class-map type inspect application [match-all] class_map_name
hostname(config-cmap)#

```

Where the *application* is the application you want to inspect. For supported applications, see [Chapter 25, “Configuring Application Layer Protocol Inspection.”](#)

The *class\_map\_name* argument is the name of the class map up to 40 characters in length.

The **match-all** keyword is the default, and specifies that traffic must match all criteria to match the class map.

The CLI enters class-map configuration mode, where you can enter one or more **match** commands.

**Step 2** (Optional) To add a description to the class map, enter the following command:

```
hostname(config-cmap)# description string
```

- Step 3** Define the traffic to include in the class by entering one or more **match** commands available for your application.

To specify traffic that should not match the class map, use the **match not** command. For example, if the **match not** command specifies the string “example.com,” then any traffic that includes “example.com” does not match the class map.

To see the **match** commands available for each application, see [Chapter 25, “Configuring Application Layer Protocol Inspection.”](#)

The following example creates an HTTP class map that must match all criteria:

```
hostname(config-cmap)# class-map type inspect http match-all http-traffic
hostname(config-cmap)# match req-resp content-type mismatch
hostname(config-cmap)# match request body length gt 1000
hostname(config-cmap)# match not request uri regex class URLs
```

## Creating a Regular Expression

A regular expression matches text strings either literally as an exact string, or by using *metacharacters* so you can match multiple variants of a text string. You can use a regular expression to match the content of certain application traffic; for example, you can match a URL string inside an HTTP packet.

Use **Ctrl+V** to escape all of the special characters in the CLI, such as question mark (?) or a tab. For example, type **d[Ctrl+V]g** to enter **d?g** in the configuration.

See the **regex** command in the *Cisco Security Appliance Command Reference* for performance impact information when matching a regular expression to packets.



### Note

As an optimization, the security appliance searches on the deobfuscated URL. Deobfuscation compresses multiple forward slashes (/) into a single slash. For strings that commonly use double slashes, like “http://”, be sure to search for “http:/" instead.

[Table 21-1](#) lists the metacharacters that have special meanings.

**Table 21-1** *regex Metacharacters*

| Character | Description   | Notes                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| .         | Dot           | Matches any single character. For example, <b>d.g</b> matches dog, dag, dtg, and any word that contains those characters, such as doggonnit.                                                                                                                                                                                                                                            |
| (exp)     | Subexpression | A subexpression segregates characters from surrounding characters, so that you can use other metacharacters on the subexpression. For example, <b>d(ola)g</b> matches dog and dag, but <b>dolag</b> matches do and ag. A subexpression can also be used with repeat quantifiers to differentiate the characters meant for repetition. For example, <b>ab(xy){3}z</b> matches abxyxyxyz. |

**Table 21-1** *regex Metacharacters (continued)*

| Character   | Description               | Notes                                                                                                                                                                                                                                                                                                                                                       |
|-------------|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|             | Alternation               | Matches either expression it separates. For example, <b>dog cat</b> matches dog or cat.                                                                                                                                                                                                                                                                     |
| ?           | Question mark             | A quantifier that indicates that there are 0 or 1 of the previous expression. For example, <b>lo?se</b> matches lse or lose.<br><br><b>Note</b> You must enter <b>Ctrl+V</b> and then the question mark or else the help function is invoked.                                                                                                               |
| *           | Asterisk                  | A quantifier that indicates that there are 0, 1 or any number of the previous expression. For example, <b>lo*se</b> matches lse, lose, loose, and so on.                                                                                                                                                                                                    |
| +           | Plus                      | A quantifier that indicates that there is at least 1 of the previous expression. For example, <b>lo+se</b> matches lose and loose, but not lse.                                                                                                                                                                                                             |
| {x}         | Repeat quantifier         | Repeat exactly <i>x</i> times. For example, <b>ab(xy){3}z</b> matches abxyxyxyz.                                                                                                                                                                                                                                                                            |
| {x,}        | Minimum repeat quantifier | Repeat at least <i>x</i> times. For example, <b>ab(xy){2,}z</b> matches abxyxyz, abxyxyxyz, and so on.                                                                                                                                                                                                                                                      |
| [abc]       | Character class           | Matches any character in the brackets. For example, <b>[abc]</b> matches a, b, or c.                                                                                                                                                                                                                                                                        |
| [^abc]      | Negated character class   | Matches a single character that is not contained within the brackets. For example, <b>[^abc]</b> matches any character other than a, b, or c. <b>[^A-Z]</b> matches any single character that is not an uppercase letter.                                                                                                                                   |
| [a-c]       | Character range class     | Matches any character in the range. <b>[a-z]</b> matches any lowercase letter. You can mix characters and ranges: <b>[abcq-z]</b> matches a, b, c, q, r, s, t, u, v, w, x, y, z, and so does <b>[a-cq-z]</b> .<br><br>The dash (-) character is literal only if it is the last or the first character within the brackets: <b>[abc-]</b> or <b>[-abc]</b> . |
| ""          | Quotation marks           | Preserves trailing or leading spaces in the string. For example, <b>" test"</b> preserves the leading space when it looks for a match.                                                                                                                                                                                                                      |
| ^           | Caret                     | Specifies the beginning of a line.                                                                                                                                                                                                                                                                                                                          |
| \           | Escape character          | When used with a metacharacter, matches a literal character. For example, <b>\[</b> matches the left square bracket.                                                                                                                                                                                                                                        |
| <i>char</i> | Character                 | When character is not a metacharacter, matches the literal character.                                                                                                                                                                                                                                                                                       |
| \r          | Carriage return           | Matches a carriage return 0x0d.                                                                                                                                                                                                                                                                                                                             |
| \n          | Newline                   | Matches a new line 0x0a.                                                                                                                                                                                                                                                                                                                                    |
| \t          | Tab                       | Matches a tab 0x09.                                                                                                                                                                                                                                                                                                                                         |
| \f          | Formfeed                  | Matches a form feed 0x0c.                                                                                                                                                                                                                                                                                                                                   |

**Table 21-1** *regex Metacharacters (continued)*

| Character | Description                | Notes                                                                                                          |
|-----------|----------------------------|----------------------------------------------------------------------------------------------------------------|
| \xNN      | Escaped hexadecimal number | Matches an ASCII character using hexadecimal (exactly two digits).                                             |
| \NNN      | Escaped octal number       | Matches an ASCII character as octal (exactly three digits). For example, the character 040 represents a space. |

To test and create a regular expression, perform the following steps:

- Step 1** To test a regular expression to make sure it matches what you think it will match, enter the following command:

```
hostname(config)# test regex input_text regular_expression
```

Where the *input\_text* argument is a string you want to match using the regular expression, up to 201 characters in length.

The *regular\_expression* argument can be up to 100 characters in length.

Use **Ctrl+V** to escape all of the special characters in the CLI. For example, to enter a tab in the input text in the **test regex** command, you must enter **test regex "test[Ctrl+V Tab]" "test\t"**.

If the regular expression matches the input text, you see the following message:

```
INFO: Regular expression match succeeded.
```

If the regular expression does not match the input text, you see the following message:

```
INFO: Regular expression match failed.
```

- Step 2** To add a regular expression after you tested it, enter the following command:

```
hostname(config)# regex name regular_expression
```

Where the *name* argument can be up to 40 characters in length.

The *regular\_expression* argument can be up to 100 characters in length.

The following example creates two regular expressions for use in an inspection policy map:

```
hostname(config)# regex url_example example\.com
hostname(config)# regex url_example2 example2\.com
```

## Creating a Regular Expression Class Map

A regular expression class map identifies one or more regular expressions. You can use a regular expression class map to match the content of certain traffic; for example, you can match URL strings inside HTTP packets.

To create a regular expression class map, perform the following steps:

- Step 1** Create one or more regular expressions according to the [“Creating a Regular Expression”](#) section.

**Step 2** Create a class map by entering the following command:

```
hostname(config)# class-map type regex match-any class_map_name
hostname(config-cmap)#
```

Where *class\_map\_name* is a string up to 40 characters in length. The name “class-default” is reserved. All types of class maps use the same name space, so you cannot reuse a name already used by another type of class map.

The **match-any** keyword specifies that the traffic matches the class map if it matches only one of the regular expressions.

The CLI enters class-map configuration mode.

**Step 3** (Optional) Add a description to the class map by entering the following command:

```
hostname(config-cmap)# description string
```

**Step 4** Identify the regular expressions you want to include by entering the following command for each regular expression:

```
hostname(config-cmap)# match regex regex_name
```

---

The following example creates two regular expressions, and adds them to a regular expression class map. Traffic matches the class map if it includes the string “example.com” or “example2.com.”

```
hostname(config)# regex url_example example\.com
hostname(config)# regex url_example2 example2\.com
hostname(config)# class-map type regex match-any URLs
hostname(config-cmap)# match regex url_example
hostname(config-cmap)# match regex url_example2
```

## Defining Actions (Layer 3/4 Policy Map)

This section describes how to associate actions with Layer 3/4 class maps by creating a Layer 3/4 policy map. This section includes the following topics:

- [Layer 3/4 Policy Map Overview, page 21-15](#)
- [Default Layer 3/4 Policy Map, page 21-18](#)
- [Adding a Layer 3/4 Policy Map, page 21-19](#)

## Layer 3/4 Policy Map Overview

This section describes how Layer 3/4 policy maps work, and includes the following topics:

- [Policy Map Guidelines, page 21-16](#)
- [Supported Feature Types, page 21-16](#)
- [Hierarchical Policy Maps, page 21-16](#)
- [Feature Directionality, page 21-17](#)
- [Feature Matching Guidelines within a Policy Map, page 21-17](#)
- [Feature Matching Guidelines for multiple Policy Maps, page 21-18](#)

- [Order in Which Multiple Feature Actions are Applied](#), page 21-18

## Policy Map Guidelines

See the following guidelines for using policy maps:

- You can only assign one policy map per interface.
- You can apply the same policy map to multiple interfaces.
- You can identify multiple Layer 3/4 class maps in a Layer 3/4 policy map.
- For each class map, you can assign multiple actions from one or more feature types.
- You can create a hierarchical policy map. See the [“Hierarchical Policy Maps”](#) section on page 21-16.

## Supported Feature Types

Feature types supported by the Modular Policy Framework that you can enable in the policy map include the following:

- QoS input policing
- TCP normalization, TCP and UDP connection limits and timeouts, and TCP sequence number randomization
- CSC
- Application inspection
- IPS
- QoS output policing
- QoS standard priority queue
- QoS traffic shaping, hierarchical priority queue

## Hierarchical Policy Maps

If you enable QoS traffic shaping for a class map, then you can optionally enable priority queueing for a subset of shaped traffic. To do so, you need to create a policy map for the priority queueing, and then within the traffic shaping policy map, you can call the priority class map. Only the traffic shaping class map is applied to an interface.

See [Chapter 24, “Configuring QoS,”](#) for more information about this feature.

Hierarchical policy maps are only supported for traffic shaping and priority queueing.

To implement a hierarchical policy map, perform the following tasks:

1. Identify the prioritized traffic according to the [“Identifying Traffic \(Layer 3/4 Class Map\)”](#) section on page 21-4.

You can create multiple class maps to be used in the hierarchical policy map.

2. Create a policy map according to the [“Defining Actions \(Layer 3/4 Policy Map\)”](#) section on page 21-15, and identify the sole action for each class map as **priority**.
3. Create a separate policy map according to the [“Defining Actions \(Layer 3/4 Policy Map\)”](#) section on page 21-15, and identify the **shape** action for the **class-default** class map.

Traffic shaping can only be applied to the **class-default** class map.

4. For the same class map, identify the priority policy map that you created in Step 2 using the **service-policy** *priority\_policy\_map* command.
5. Apply the shaping policy map to the interface according to [“Applying Actions to an Interface \(Service Policy\)” section on page 21-21](#).

## Feature Directionality

Actions are applied to traffic bidirectionally or unidirectionally depending on the feature. For features that are applied bidirectionally, all traffic that enters or exits the interface to which you apply the policy map is affected if the traffic matches the class map for both directions.



### Note

When you use a global policy, all features are unidirectional; features that are normally bidirectional when applied to a single interface only apply to the ingress of each interface when applied globally. Because the policy is applied to all interfaces, the policy will be applied in both directions so bidirectionality in this case is redundant.

For features that are applied unidirectionally, for example QoS priority queue, only traffic that exits the interface to which you apply the policy map is affected. See [Table 21-2](#) for the directionality of each feature.

**Table 21-2 Feature Directionality**

| Feature                                                                                              | Single Interface Direction | Global Direction |
|------------------------------------------------------------------------------------------------------|----------------------------|------------------|
| TCP normalization, TCP and UDP connection limits and timeouts, and TCP sequence number randomization | Bidirectional              | Ingress          |
| CSC                                                                                                  | Bidirectional              | Ingress          |
| Application inspection                                                                               | Bidirectional              | Ingress          |
| IPS                                                                                                  | Bidirectional              | Ingress          |
| QoS input policing                                                                                   | Ingress                    | Ingress          |
| QoS output policing                                                                                  | Egress                     | Egress           |
| QoS standard priority queue                                                                          | Egress                     | Egress           |
| QoS traffic shaping, hierarchical priority queue                                                     | Egress                     | Egress           |

## Feature Matching Guidelines within a Policy Map

See the following guidelines for how a packet matches class maps in a policy map:

- A packet can match only one class map in the policy map for each feature type.
- When the packet matches a class map for a feature type, the security appliance does not attempt to match it to any subsequent class maps for that feature type.
- If the packet matches a subsequent class map for a different feature type, however, then the security appliance also applies the actions for the subsequent class map.

For example, if a packet matches a class map for connection limits, and also matches a class map for application inspection, then both class map actions are applied.

If a packet matches a class map for application inspection, but also matches another class map that includes application inspection, then the second class map actions are not applied.

## Feature Matching Guidelines for multiple Policy Maps

For TCP and UDP traffic (and ICMP when you enable stateful ICMP inspection), Modular Policy Framework operates on traffic flows, and not just individual packets. If traffic is part of an existing connection that matches a feature in a policy on one interface, that traffic flow cannot also match the same feature in a policy on another interface; only the first policy is used.

For example, if HTTP traffic matches a policy on the inside interface to inspect HTTP traffic, and you have a separate policy on the outside interface for HTTP inspection, then that traffic is not also inspected on the egress of the outside interface. Similarly, the return traffic for that connection will not be inspected by the ingress policy of the outside interface, nor by the egress policy of the inside interface.

For traffic that is not treated as a flow, for example ICMP when you do not enable stateful ICMP inspection, returning traffic can match a different policy map on the returning interface. For example, if you configure IPS inspection on the inside and outside interfaces, but the inside policy uses virtual sensor 1 while the outside policy uses virtual sensor 2, then a non-stateful Ping will match virtual sensor 1 outbound, but will match virtual sensor 2 inbound.

## Order in Which Multiple Feature Actions are Applied

The order in which different types of actions in a policy map are performed is independent of the order in which the actions appear in the policy map. Actions are performed in the following order:

- QoS input policing
- TCP normalization, TCP and UDP connection limits and timeouts, and TCP sequence number randomization



**Note** When a the security appliance performs a proxy service (such as AAA or CSC) or it modifies the TCP payload (such as FTP inspection), the TCP normalizer acts in dual mode, where it is applied before and after the proxy or payload modifying service.

- CSC
- Application inspection
- IPS
- QoS output policing
- QoS standard priority queue
- QoS traffic shaping, hierarchical priority queue

## Default Layer 3/4 Policy Map

The configuration includes a default Layer 3/4 policy map that the security appliance uses in the default global policy. It is called **global\_policy** and performs inspection on the default inspection traffic. You can only apply one global policy, so if you want to alter the global policy, you need to either reconfigure the default policy or disable it and apply a new one.

The default policy map configuration includes the following commands:



```

policy-map global_policy
class inspection_default
inspect dns preset_dns_map
inspect ftp
inspect h323 h225
inspect h323 ras
inspect rsh
inspect rtsp
inspect esmtp
inspect sqlnet
inspect skinny
inspect sunrpc
inspect xdmcp
inspect sip
inspect netbios
inspect tftp

```

## Adding a Layer 3/4 Policy Map

The maximum number of policy maps is 64. To create a Layer 3/4 policy map, perform the following steps:

**Step 1** Add the policy map by entering the following command:

```
hostname(config)# policy-map policy_map_name
```

The *policy\_map\_name* argument is the name of the policy map up to 40 characters in length. All types of policy maps use the same name space, so you cannot reuse a name already used by another type of policy map. The CLI enters policy-map configuration mode.

**Step 2** (Optional) Specify a description for the policy map:

```
hostname(config-pmap)# description text
```

**Step 3** Specify a previously configured Layer 3/4 class map using the following command:

```
hostname(config-pmap)# class class_map_name
```

See the “[Identifying Traffic \(Layer 3/4 Class Map\)](#)” section on page 21-4 to add a class map.

**Step 4** Specify one or more actions for this class map.

- IPS. See the “[Diverting Traffic to the AIP SSM](#)” section on page 22-2.
- CSC. See the “[Diverting Traffic to the CSC SSM](#)” section on page 22-11.
- TCP normalization. See the “[Configuring TCP Normalization](#)” section on page 23-1.
- TCP and UDP connection limits and timeouts, and TCP sequence number randomization. See the “[Configuring Connection Limits and Timeouts](#)” section on page 23-6.
- QoS. See [Chapter 24, “Configuring QoS.”](#)



**Note**

You can configure a hierarchical policy map for the traffic shaping and priority queue features. See the “[Hierarchical Policy Maps](#)” section on page 21-16 for more information.

- Application inspection. See [Chapter 25, “Configuring Application Layer Protocol Inspection.”](#)

**Note**

If there is no **match default\_inspection\_traffic** command in a class map, then at most one **inspect** command is allowed to be configured under the class.

**Step 5** Repeat [Step 3](#) and [Step 4](#) for each class map you want to include in this policy map.

The following is an example of a **policy-map** command for connection policy. It limits the number of connections allowed to the web server 10.1.1.1:

```
hostname(config)# access-list http-server permit tcp any host 10.1.1.1
hostname(config)# class-map http-server
hostname(config-cmap)# match access-list http-server

hostname(config)# policy-map global-policy
hostname(config-pmap)# description This policy map defines a policy concerning connection
to http server.
hostname(config-pmap)# class http-server
hostname(config-pmap-c)# set connection conn-max 256
```

The following example shows how multi-match works in a policy map:

```
hostname(config)# class-map inspection_default
hostname(config-cmap)# match default-inspection-traffic
hostname(config)# class-map http_traffic
hostname(config-cmap)# match port tcp eq 80

hostname(config)# policy-map outside_policy
hostname(config-pmap)# class inspection_default
hostname(config-pmap-c)# inspect http http_map
hostname(config-pmap-c)# inspect sip
hostname(config-pmap)# class http_traffic
hostname(config-pmap-c)# set connection timeout tcp 0:10:0
```

The following example shows how traffic matches the first available class map, and will not match any subsequent class maps that specify actions in the same feature domain:

```
hostname(config)# class-map telnet_traffic
hostname(config-cmap)# match port tcp eq 23
hostname(config)# class-map ftp_traffic
hostname(config-cmap)# match port tcp eq 21
hostname(config)# class-map tcp_traffic
hostname(config-cmap)# match port tcp range 1 65535
hostname(config)# class-map udp_traffic
hostname(config-cmap)# match port udp range 0 65535
hostname(config)# policy-map global_policy
hostname(config-pmap)# class telnet_traffic
hostname(config-pmap-c)# set connection timeout tcp 0:0:0
hostname(config-pmap-c)# set connection conn-max 100
hostname(config-pmap)# class ftp_traffic
hostname(config-pmap-c)# set connection timeout tcp 0:5:0
hostname(config-pmap-c)# set connection conn-max 50
hostname(config-pmap)# class tcp_traffic
hostname(config-pmap-c)# set connection timeout tcp 2:0:0
hostname(config-pmap-c)# set connection conn-max 2000
```

When a Telnet connection is initiated, it matches **class telnet\_traffic**. Similarly, if an FTP connection is initiated, it matches **class ftp\_traffic**. For any TCP connection other than Telnet and FTP, it will match **class tcp\_traffic**. Even though a Telnet or FTP connection can match **class tcp\_traffic**, the security appliance does not make this match because they previously matched other classes.

## Applying Actions to an Interface (Service Policy)

To activate the Layer 3/4 policy map, create a service policy that applies it to one or more interfaces or that applies it globally to all interfaces. Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with inspections, and an interface policy with TCP normalization, then both inspections and TCP normalization are applied to the interface. However, if you have a global policy with inspections, and an interface policy with inspections, then only the interface policy inspections are applied to that interface.

- To create a service policy by associating a policy map with an interface, enter the following command:

```
hostname(config)# service-policy policy_map_name interface interface_name
```

- To create a service policy that applies to all interfaces that do not have a specific policy, enter the following command:

```
hostname(config)# service-policy policy_map_name global
```

By default, the configuration includes a global policy that matches all default application inspection traffic and applies inspection to the traffic globally. You can only apply one global policy, so if you want to alter the global policy, you need to either edit the default policy or disable it and apply a new one.

The default service policy includes the following command:

```
service-policy global_policy global
```

For example, the following command enables the inbound\_policy policy map on the outside interface:

```
hostname(config)# service-policy inbound_policy interface outside
```

The following commands disable the default global policy, and enables a new one called new\_global\_policy on all other security appliance interfaces:

```
hostname(config)# no service-policy global_policy global
hostname(config)# service-policy new_global_policy global
```

## Modular Policy Framework Examples


This section includes several Modular Policy Framework examples, and includes the following topics:

- [Applying Inspection and QoS Policing to HTTP Traffic, page 21-22](#)
- [Applying Inspection to HTTP Traffic Globally, page 21-22](#)
- [Applying Inspection and Connection Limits to HTTP Traffic to Specific Servers, page 21-23](#)
- [Applying Inspection to HTTP Traffic with NAT, page 21-24](#)

## Applying Inspection and QoS Policing to HTTP Traffic

In this example (see [Figure 21-1](#)), any HTTP connection (TCP traffic on port 80) that enters or exits the security appliance through the outside interface is classified for HTTP inspection. Any HTTP traffic that exits the outside interface is classified for policing.

**Figure 21-1** HTTP Inspection and QoS Policing



See the following commands for this example:


```
hostname(config)# class-map http_traffic
hostname(config-cmap)# match port tcp eq 80

hostname(config)# policy-map http_traffic_policy
hostname(config-pmap)# class http_traffic
hostname(config-pmap-c)# inspect http
hostname(config-pmap-c)# police output 250000
hostname(config)# service-policy http_traffic_policy interface outside
```

## Applying Inspection to HTTP Traffic Globally

In this example (see [Figure 21-2](#)), any HTTP connection (TCP traffic on port 80) that enters the security appliance through any interface is classified for HTTP inspection. Because the policy is a global policy, inspection occurs only as the traffic enters each interface.

**Figure 21-2** Global HTTP Inspection



See the following commands for this example:

```
hostname(config)# class-map http_traffic
hostname(config-cmap)# match port tcp eq 80
```

```

hostname(config)# policy-map http_traffic_policy
hostname(config-pmap)# class http_traffic
hostname(config-pmap-c)# inspect http
hostname(config)# service-policy http_traffic_policy global


```

## Applying Inspection and Connection Limits to HTTP Traffic to Specific Servers

In this example (see Figure 21-3), any HTTP connection destined for Server A (TCP traffic on port 80) that enters the security appliance through the outside interface is classified for HTTP inspection and maximum connection limits. Connections initiated from server A to Host A does not match the access list in the class map, so it is not affected.

Any HTTP connection destined for Server B that enters the security appliance through the inside interface is classified for HTTP inspection. Connections initiated from server B to Host B does not match the access list in the class map, so it is not affected.

**Figure 21-3** HTTP Inspection and Connection Limits to Specific Servers



See the following commands for this example:

```

hostname(config)# static (inside,outside) 209.165.201.1 192.168.1.2
hostname(config)# nat (inside) 1 192.168.1.0 255.255.255.0
hostname(config)# global (outside) 1 209.165.201.2
hostname(config)# access-list serverA extended permit tcp any host 209.165.201.1 eq 80
hostname(config)# access-list ServerB extended permit tcp any host 209.165.200.227 eq 80

hostname(config)# class-map http_serverA
hostname(config-cmap)# match access-list serverA
hostname(config)# class-map http_serverB
hostname(config-cmap)# match access-list serverB

hostname(config)# policy-map policy_serverA
hostname(config-pmap)# class http_serverA
hostname(config-pmap-c)# inspect http
hostname(config-pmap-c)# set connection conn-max 100
hostname(config)# policy-map policy_serverB
hostname(config-pmap)# class http_serverB
hostname(config-pmap-c)# inspect http


hostname(config)# service-policy policy_serverB interface inside
hostname(config)# service-policy policy_serverA interface outside

```

## Applying Inspection to HTTP Traffic with NAT

In this example, the Host on the inside network has two addresses: one is the real IP address 192.168.1.1, and the other is a mapped IP address used on the outside network, 209.165.200.225. Because the policy is applied to the inside interface, where the real address is used, then you must use the real IP address in the access list in the class map. If you applied it to the outside interface, you would use the mapped address.

**Figure 21-4 HTTP Inspection with NAT**



See the following commands for this example:

```
hostname(config)# static (inside,outside) 209.165.200.225 192.168.1.1
hostname(config)# access-list http_client extended permit tcp host 192.168.1.1 any eq 80

hostname(config)# class-map http_client
hostname(config-cmap)# match access-list http_client

hostname(config)# policy-map http_client
hostname(config-pmap)# class http_client
hostname(config-pmap-c)# inspect http

hostname(config)# service-policy http_client interface inside
```



## CHAPTER 22

# Managing AIP SSM and CSC SSM

---

The Cisco ASA 5500 series adaptive security appliance supports a variety of SSMs. This chapter describes how to configure the adaptive security appliance to support an AIP SSM or a CSC SSM, including how to send traffic to these SSMs.

For information about the 4GE SSM for the ASA 5000 series adaptive security appliance, see [Chapter 5, “Configuring Ethernet Settings and Subinterfaces”](#).



### Note

---

The Cisco PIX 500 series security appliances cannot support SSMs.

---

This chapter includes the following sections:

- [Managing the AIP SSM, page 22-1](#)
- [Managing the CSC SSM, page 22-5](#)
- [Checking SSM Status, page 22-13](#)
- [Transferring an Image onto an SSM, page 22-14](#)

## Managing the AIP SSM

This section contains the following topics:

- [About the AIP SSM, page 22-1](#)
- [Getting Started with the AIP SSM, page 22-2](#)
- [Diverting Traffic to the AIP SSM, page 22-2](#)
- [Sessioning to the AIP SSM and Running Setup, page 22-4](#)

## About the AIP SSM

The ASA 5500 series adaptive security appliance supports the AIP SSM, which runs advanced IPS software that provides further security inspection. The adaptive security appliance diverts packets to the AIP SSM just before the packet exits the egress interface (or before VPN encryption occurs, if configured) and after other firewall policies are applied. For example, packets that are blocked by an access list are not forwarded to the AIP SSM.

The AIP SSM can operate in one of two modes, as follows:

- **Inline mode**—Places the AIP SSM directly in the traffic flow. No traffic can continue through the adaptive security appliance without first passing through, and being inspected by, the AIP SSM. This mode is the most secure because every packet is analyzed before being allowed through. Also, the AIP SSM can implement a blocking policy on a packet-by-packet basis. This mode, however, can affect throughput. You specify this mode with the **inline** keyword of the **ips** command.
- **Promiscuous mode**—Sends a duplicate stream of traffic to the AIP SSM. This mode is less secure, but has little impact on traffic throughput. Unlike operation in inline mode, the SSM operating in promiscuous mode can only block traffic by instructing the adaptive security appliance to **shun** the traffic or by resetting a connection on the adaptive security appliance. Also, while the AIP SSM is analyzing the traffic, a small amount of traffic might pass through the adaptive security appliance before the AIP SSM can block it. You specify this mode with the **inline** keyword of the **ips** command.

You can specify how the adaptive security appliance treats traffic when the AIP SSM is unavailable due to hardware failure or other causes. Two keywords of the **ips** command control this behavior. The **fail-close** keyword sets the adaptive security appliance to block all traffic if the AIP SSM is unavailable. The **fail-open** keyword sets the adaptive security appliance to allow all traffic through, uninspected, if the AIP SSM is unavailable.

For more information about configuring the operating mode of the AIP SSM and how the adaptive security appliance treats traffic during an AIP SSM failure, see the [“Diverting Traffic to the AIP SSM” section on page 22-2](#).

## Getting Started with the AIP SSM

Configuring the AIP SSM is a two-part process that involves configuration of the ASA 5500 series adaptive security appliance first, and then configuration of the AIP SSM:

1. On the ASA 5500 series adaptive security appliance, identify traffic to divert to the AIP SSM (as described in the [“Diverting Traffic to the AIP SSM” section on page 22-2](#)).
2. On the AIP SSM, configure the inspection and protection policy, which determines how to inspect traffic and what to do when an intrusion is detected. Because the IPS software that runs on the AIP SSM is very robust and beyond the scope of this document, detailed configuration information is available in the following separate documentation:
  - [Configuring the Cisco Intrusion Prevention System Sensor Using the Command Line Interface](#)
  - [Cisco Intrusion Prevention System Command Reference](#)

## Diverting Traffic to the AIP SSM

You use MPF commands to configure the adaptive security appliance to divert traffic to the AIP SSM. Before configuring the adaptive security appliance to do so, read [Chapter 21, “Using Modular Policy Framework,”](#) which introduces MPF concepts and common commands.

To identify traffic to divert from the adaptive security appliance to the AIP SSM, perform the following steps:

---

**Step 1** Create an access list that matches all traffic:

```
hostname(config)# access-list acl-name permit ip any any
```



- Step 2** Create a class map to identify the traffic that should be diverted to the AIP SSM. Use the **class-map** command to do so, as follows:

```
hostname(config)# class-map class_map_name
hostname(config-cmap)#
```

where *class\_map\_name* is the name of the traffic class. When you enter the **class-map** command, the CLI enters class map configuration mode.

- Step 3** With the access list you created in [Step 1](#), use a **match access-list** command to identify the traffic to be scanned:

```
hostname(config-cmap)# match access-list acl-name
```

- Step 4** Create a policy map or modify an existing policy map that you want to use to send traffic to the AIP SSM. To do so, use the **policy-map** command, as follows.

```
hostname(config-cmap)# policy-map policy_map_name
hostname(config-pmap)#
```

where *policy\_map\_name* is the name of the policy map. The CLI enters the policy map configuration mode and the prompt changes accordingly.

- Step 5** Specify the class map, created in [Step 2](#), that identifies the traffic to be scanned. Use the **class** command to do so, as follows.

```
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

where *class\_map\_name* is the name of the class map you created in [Step 2](#). The CLI enters the policy map class configuration mode and the prompt changes accordingly.

- Step 6** Assign the traffic identified by the class map as traffic to be sent to the AIP SSM. Use the **ips** command to do so, as follows.

```
hostname(config-pmap-c)# ips {inline | promiscuous} {fail-close | fail-open}
```

The **inline** and **promiscuous** keywords control the operating mode of the AIP SSM. The **fail-close** and **fail-open** keywords control how the adaptive security appliance treats traffic when the AIP SSM is unavailable. For more information about the operating modes and failure behavior, see the [“About the AIP SSM”](#) section on page 22-1.

- Step 7** Use the **service-policy** command to apply the policy map globally or to a specific interface, as follows:

```
hostname(config-pmap-c)# service-policy policy_map_name [global | interface interface_ID]
hostname(config)#
```

where *policy\_map\_name* is the policy map you configured in [Step 4](#). If you want to apply the policy map to traffic on all the interfaces, use the **global** keyword. If you want to apply the policy map to traffic on a specific interface, use the **interface interface\_ID** option, where *interface\_ID* is the name assigned to the interface with the **nameif** command.

Only one global policy is allowed. You can override the global policy on an interface by applying a service policy to that interface. You can only apply one policy map to each interface.

The adaptive security appliance begins diverting traffic to the AIP SSM as specified.

---

The following example diverts all IP traffic to the AIP SSM in promiscuous mode, and blocks all IP traffic should the AIP SSM card fail for any reason:

```
hostname(config)# access-list IPS permit ip any any
hostname(config)# class-map my-ips-class
```

```
hostname(config-cmap)# match access-list IPS
hostname(config-cmap)# policy-map my-ips-policy
hostname(config-pmap)# class my-ips-class
hostname(config-pmap-c)# ips promiscuous fail-close
hostname(config-pmap-c)# service-policy my-ips-policy global
```

For a complete example of network traffic diversion from the adaptive security appliance to the AIP SSM, see [Example 16: Network Traffic Diversion](#).

## Sessioning to the AIP SSM and Running Setup

After you have completed configuration of the ASA 5500 series adaptive security appliance to divert traffic to the AIP SSM, session to the AIP SSM and run the setup utility for initial configuration.



### Note

You can either session to the SSM from the adaptive security appliance (by using the **session 1** command) or you can connect directly to the SSM using SSH or Telnet on its management interface. Alternatively, you can use ASDM.

To session to the AIP SSM from the adaptive security appliance, perform the following steps:

- Step 1** Enter the **session 1** command to session from the ASA 5500 series adaptive security appliance to the AIP SSM:

```
hostname# session 1
Opening command session with slot 1.
Connected to slot 1. Escape character sequence is 'CTRL-^X'.
```

- Step 2** Enter the username and password. The default username and password are both **cisco**.



### Note

The first time you log in to the AIP SSM you are prompted to change the default password. Passwords must be at least eight characters long and *not* a dictionary word.

```
login: cisco
Password:
Last login: Fri Sep 2 06:21:20 from xxx.xxx.xxx.xxx
NOTICE
This product contains cryptographic features and is subject to United States
and local country laws governing import, export, transfer and use. Delivery
of Cisco cryptographic products does not imply third-party authority to import,
export, distribute or use encryption. Importers, exporters, distributors and
users are responsible for compliance with U.S. and local country laws. By using
this product you agree to comply with applicable laws and regulations. If you
are unable to comply with U.S. and local laws, return this product immediately.
```

A summary of U.S. laws governing Cisco cryptographic products may be found at:  
<http://www.cisco.com/wwl/export/crypto/tool/stqrg.html>

If you require further assistance please contact us by sending email to  
[export@cisco.com](mailto:export@cisco.com).

```
LICENSE NOTICE
There is no license key installed on the system.
Please go to http://www.cisco.com/go/license
to obtain a new license or install a license.
AIP SSM#
```

**Note**

If you see the preceding license notice (which displays only in some versions of software), you can ignore the message until you need to upgrade the signature files on the AIP SSM. The AIP SSM continues to operate at the current signature level until a valid license key is installed. You can install the license key at a later time. The license key does not affect the current functionality of the AIP SSM.

**Step 3** Enter the **setup** command to run the setup utility for initial configuration of the AIP SSM:

```
AIP SSM# setup
```

You are now ready to configure the AIP SSM for intrusion prevention. See the following two guides for AIP SSM configuration information:

- [Configuring the Cisco Intrusion Prevention System Sensor Using the Command Line Interface](#)
- [Cisco Intrusion Prevention System Command Reference](#)

## Managing the CSC SSM

This section contains the following topics:

- [About the CSC SSM, page 22-5](#)
- [Getting Started with the CSC SSM, page 22-7](#)
- [Determining What Traffic to Scan, page 22-9](#)
- [Limiting Connections Through the CSC SSM, page 22-11](#)
- [Diverting Traffic to the CSC SSM, page 22-11](#)

## About the CSC SSM

The ASA 5500 series adaptive security appliance supports the CSC SSM, which runs Content Security and Control software. The CSC SSM provides protection against viruses, spyware, spam, and other unwanted traffic. It accomplishes this by scanning the FTP, HTTP, POP3, and SMTP traffic that you configure the adaptive security appliance to send to it.


[Figure 22-1](#) illustrates the flow of traffic through an adaptive security appliance that has the following:

- A CSC SSM installed and setup.
- A service policy that determines what traffic is diverted to the SSM for scans.

In this example, the client could be a network user who is accessing a website, downloading files from an FTP server, or retrieving mail from a POP3 server. SMTP scans differ in that you should configure the adaptive security appliance to scan traffic sent from outside to SMTP servers protected by the adaptive security appliance.

**Note**

The CSC SSM can scan FTP file transfers only when FTP inspection is enabled on the adaptive security appliance. By default, FTP inspection is enabled.

**Figure 22-1** Flow of Scanned Traffic with CSC SSM

You use ASDM for system setup and monitoring of the CSC SSM. For advanced configuration of content security policies in the CSC SSM software, you access the web-based GUI for the CSC SSM by clicking links within ASDM. Use of the CSC SSM GUI is explained in the *Trend Micro InterScan for Cisco CSC SSM Administrator Guide*.


**Note**

ASDM and the CSC SSM maintain separate passwords. You can configure their passwords to be identical; however, changing one of these two passwords does not affect the other password.

The connection between the host running ASDM and the adaptive security appliance is made through a management port on the adaptive security appliance. The connection to the CSC SSM GUI is made through the SSM management port. Because these two connections are required to manage the CSC SSM, any host running ASDM must be able to reach the IP address of both the adaptive security appliance management port and the SSM management port.

Figure 22-2 shows an adaptive security appliance with a CSC SSM that is connected to a dedicated management network. While use of a dedicated management network is not required, we recommend it. Of particular interest in Figure 22-2 are the following:

- An HTTP proxy server is connected to the inside network and to the management network. This enables the CSC SSM to contact the Trend Micro update server.
- The management port of the adaptive security appliance is connected to the management network. To permit management of the adaptive security appliance and the CSC SSM, hosts running ASDM must be connected to the management network.
- The management network includes an SMTP server for email notifications for the CSC SSM and a syslog server that the CSC SSM can send syslog messages to.

**Figure 22-2** CSC SSM Deployment with a Management Network

CSC SSM cannot support stateful failover, because the CSC SSM does not maintain connection information and therefore cannot provide the failover unit with information necessary for stateful failover. The connections that a CSC SSM is scanning are dropped upon failure of the security appliance that the CSC SSM is installed in. When the standby adaptive security appliance becomes active, it will forward the scanned traffic to its CSC SSM and the connections will be reset.

## Getting Started with the CSC SSM

Before you receive the security benefits provided by a CSC SSM, you must perform several steps beyond simple hardware installation of the SSM. This procedure provides an overview of those steps.

To configure the adaptive security appliance and the CSC SSM, follow these steps:

- Step 1** If the CSC SSM did not come pre-installed in a Cisco ASA 5500 series adaptive security appliance, install it and connect a network cable to the management port of the SSM. For assistance with installation and connecting the SSM, see the *Cisco ASA 5500 Series Hardware Installation Guide*.

The management port of the CSC SSM must be connected to your network to allow management of and automatic updates to the CSC SSM software. Additionally, the CSC SSM uses the management port for email notifications and syslogging.

- Step 2** With the CSC SSM, you should have received a Product Authorization Key (PAK). Use the PAK to register the CSC SSM at the following URL.

<http://www.cisco.com/go/license>

After you register, you will receive activation keys by email. The activation keys are required before you can complete [Step 6](#)

- Step 3** Gather the following information, for use in [Step 6](#).

- Activation keys, received after completing [Step 2](#).
- SSM management port IP address, netmask, and gateway IP address.

**Note**

The SSM management port IP address must be accessible by the hosts used to run ASDM. The IP addresses for the SSM management port and the adaptive security appliance management interface can be in different subnets.

- DNS server IP address.
- HTTP proxy server IP address (required only if your security policies require use of a proxy server for HTTP access to the Internet).
- Domain name and hostname for the SSM.
- An email address and an SMTP server IP address and port number, for email notifications.
- IP addresses of hosts or networks allowed to manage the CSC SSM.
- Password for the CSC SSM.

**Step 4** In a web browser, access ASDM for the adaptive security appliance that the CSC SSM is in.

**Note**

If you are accessing ASDM for the first time, see the *Cisco ASA 5500 Series Adaptive Security Appliance Getting Started Guide* for assistance with the Startup Wizard.

For more information about enabling ASDM access, see the [“Allowing HTTPS Access for ASDM” section on page 40-3](#).

**Step 5** Verify time settings on the adaptive security appliance. Time setting accuracy is important for logging of security events and for automatic updates of CSC SSM software.

- If you manually control time settings, verify the clock settings, including time zone. Choose **Configuration > Properties > Device Administration > Clock**.
- If you are using NTP, verify the NTP configuration. Choose **Configuration > Properties > Device Administration > NTP**.

**Step 6** In ASDM, run the Content Security setup wizard. To do so, access the ASDM GUI in a supported web browser and on the Home page, click the **Content Security** tab. The Content Security setup wizard runs. For assistance with the Content Security setup wizard, click the **Help** button.

**Note**

If you are accessing ASDM for the first time, see the *Cisco ASA 5500 Series Adaptive Security Appliance Getting Started Guide* for assistance with the Startup Wizard.

**Step 7** On the ASA 5500 series adaptive security appliance, identify traffic to divert to the CSC SSM (as described in the [“Diverting Traffic to the CSC SSM” section on page 22-11](#)).

**Step 8** (Optional) Review the default content security policies in the CSC SSM GUI. The default content security policies are suitable for most implementations. Modifying them is advanced configuration that you should perform only after reading the *Trend Micro InterScan for Cisco CSC SSM Administrator Guide*.

You review the content security policies by viewing the enabled features in the CSC SSM GUI. The availability of features depends on the license level you purchased. By default, all features included in the license you purchased are enabled.

With a Base License, the features enabled by default are SMTP virus scanning, POP3 virus scanning and content filtering, webmail virus scanning, HTTP file blocking, FTP virus scanning and file blocking, logging, and automatic updates.

With a Plus License, the additional features enabled by default are SMTP anti-spam, SMTP content filtering, POP3 anti-spam, URL blocking, and URL filtering.

To access the CSC SSM GUI, in ASDM choose **Configuration > Trend Micro Content Security**, and then select one of the following: **Web**, **Mail**, **File Transfer**, or **Updates**. The blue links on these panes, beginning with the word “Configure”, open the CSC SSM GUI.

---

## Determining What Traffic to Scan

The CSC SSM can scan FTP, HTTP, POP3, and SMTP traffic. It supports these protocols only when the destination port of the packet requesting the connection is the well known port for the protocol, that is, CSC SSM can scan only the following connections:

- FTP connections opened to TCP port 21.
- HTTP connections opened to TCP port 80.
- POP3 connections opened to TCP port 110.
- SMTP connections opened to TCP port 25.

You can choose to scan traffic for all of these protocols or any combination of them. For example, if you do not allow network users to receive POP3 email, you would not want to configure the adaptive security appliance to divert POP3 traffic to the CSC SSM (you would want to block it instead).

To maximize performance of the adaptive security appliance and the CSC SSM, divert to the CSC SSM only the traffic that you want the CSC SSM to scan. Needlessly diverting traffic that you do not want to scan, such as traffic between a trusted source and destination, can adversely affect network performance.

The action of scanning traffic with the CSC SSM is enabled with the **csc** command, which must be part of a service policy. Service policies can be applied globally or to specific interfaces; therefore, you can choose to enable the **csc** command globally or for specific interfaces.

Adding the **csc** command to your global policy ensures that all unencrypted connections through the adaptive security appliance are scanned by the CSC SSM; however, this may mean that traffic from trusted sources is needlessly scanned.

If you enable the **csc** command in interface-specific service policies, it is bi-directional. This means that when the adaptive security appliance opens a new connection, if the **csc** command is active on either the inbound or the outbound interface of the connection and if the class map for the policy identifies traffic for scanning, the adaptive security appliance diverts it to the CSC SSM.

However, bi-directionality means that if you divert to the CSC SSM any of the supported traffic types that cross a given interface, the CSC SSM is likely performing needless scans on traffic from your trusted inside networks. For example, URLs and files requested from web servers on a DMZ network are unlikely to pose content security risks to hosts on an inside network and you probably do not want the adaptive security appliance to divert such traffic to the CSC SSM.


Therefore, we highly recommend using access lists to further limit the traffic selected by the class maps of CSC SSM service policies. Specifically, use access lists that match the following:

- HTTP connections to outside networks.
- FTP connections from clients inside the adaptive security appliance to servers outside the adaptive security appliance.
- POP3 connections from clients inside the security appliance to servers outside the adaptive security appliance.

- Incoming SMTP connections destined to inside mail servers.

In Figure 22-3, the adaptive security appliance should be configured to divert traffic to CSC SSM requests from clients on the inside network for HTTP, FTP, and POP3 connections to the outside network and incoming SMTP connections from outside hosts to the mail server on the DMZ network. HTTP requests from the inside network to the web server on the DMZ network should not be scanned.

**Figure 22-3 Common Network Configuration for CSC SSM Scanning**



There are many ways you could configure the adaptive security appliance to identify the traffic that you want to scan. One approach is to define two service policies, one on the inside interface and the other on the outside interface, each with an access list that matches traffic to be scanned. The following access list could be used on the policy applied to the inside interface:

```
access-list csc_out permit tcp 192.168.10.0 255.255.255.0 any eq 21
access-list csc_out deny tcp 192.168.10.0 255.255.255.0 192.168.20.0 255.255.255.0 eq 80
access-list csc_out permit tcp 192.168.10.0 255.255.255.0 any eq 80
access-list csc_out permit tcp 192.168.10.0 255.255.255.0 any eq 110
```

As previously mentioned, policies applying the **csc** command to a specific interface are effective on both ingress and egress traffic, but by specifying 192.168.10.0 as the source network in the **csc\_out** access list the policy applied to the inside interface matches only connections initiated by the hosts on the inside network. Notice also that the second ACE of the access list uses the **deny** keyword. This ACE does **not** mean the adaptive security appliance blocks traffic sent from the 192.168.10.0 network to TCP port 80 on the 192.168.20.0 network. It simply exempts the traffic from being matched by the policy map and thus prevents the adaptive security appliance from sending it to the CSC SSM.

You can use deny statements in an access list to exempt connections with trusted external hosts from being scanned. For example, to reduce the load on the CSC SSM, you might want to exempt HTTP traffic to a well known, trusted site. If the web server at such a site had the IP address 209.165.201.7, you could add the following ACE to the **csc\_out** access list to exclude HTTP connections between the trusted external web server and inside hosts from being scanned by CSC SSM:

```
access-list csc_out deny tcp 192.168.10.0 255.255.255.0 209.165.201.7 255.255.255.255 eq 80
```

The second policy in this example, applied to the outside interface, could use the following access list:

```
access-list csc_in permit tcp any 192.168.20.0 255.255.255.0 eq 25
```



This access list matches inbound SMTP connections from any external host to any host on the DMZ network. The policy applied to the outside interface would therefore ensure that incoming SMTP email would be diverted to the CSC SSM for scanning. It would not match SMTP connections from hosts on the inside network to the mail server on the DMZ network because those connections never use the outside interface.

If the web server on the DMZ network receives files uploaded by HTTP from external hosts, you could add the following ACE to the `csc_in` access list to use the CSC SSM to protect the web server from infected files:

```
access-list csc_in permit tcp any 192.168.20.0 255.255.255.0 eq 80
```

For a complete example service policy configuration using the access lists in this section, see [Example 22-1](#).

## Limiting Connections Through the CSC SSM

The adaptive security appliance can prevent the CSC SSM and the destinations of connections it scans from accepting or even receiving requests for more connections than desired. It can do so for embryonic connections or fully established connections. Also, you can specify limits for all clients included in a class-map and per-client limits. The **set connection** command lets you configure limits for embryonic connections or fully established connections.

Also, you can specify limits for all clients included in a class-map and per-client limits. The **per-client-embryonic-max** and **per-client-max** parameters limit the maximum number of connections that individual clients can open. If a client uses more network resources simultaneously than is desired, you can use these parameters to limit the number of connections that the adaptive security appliance allows each client.

DoS attacks seek to disrupt networks by overwhelming the capacity of key hosts with connections or requests for connections. You can use the **set connection** command to thwart DoS attacks. After you configure a per-client maximum that can be supported by hosts likely to be attacked, malicious clients will be unable to overwhelm hosts on protected networks.

Use of the **set connection** command to protect the CSC SSM and the destinations of connections it scans is included in the [“Diverting Traffic to the CSC SSM”](#) section on page 22-11.

## Diverting Traffic to the CSC SSM

You use MPF commands to configure the adaptive security appliance to divert traffic to the CSC SSM. Before configuring the adaptive security appliance to do so, read [Chapter 21, “Using Modular Policy Framework,”](#) which introduces MPF concepts and common commands.

To identify traffic to divert from the adaptive security appliance to the CSC SSM, perform the following steps:

- 
- Step 1** Create an access list that matches the traffic you want scanned by the CSC SSM. To do so, use the **access-list extended** command. Create as many ACEs as needed to match all the traffic. For example, if you want to specify FTP, HTTP, POP3, and SMTP traffic, you would need four ACEs. For guidance on identifying the traffic you want to scan, see the [“Determining What Traffic to Scan”](#) section on page 22-9.
  - Step 2** Create a class map to identify the traffic that should be diverted to the CSC SSM. Use the **class-map** command to do so, as follows.

```
hostname(config)# class-map class_map_name
hostname(config-cmap)#
```

where *class\_map\_name* is the name of the traffic class. When you enter the **class-map** command, the CLI enters class map configuration mode.

- Step 3** With the access list you created in [Step 1](#), use a **match access-list** command to identify the traffic to be scanned:

```
hostname(config-cmap)# match access-list acl-name
```

- Step 4** Create a policy map or modify an existing policy map that you want to use to send traffic to the CSC SSM. To do so, use the **policy-map** command, as follows.

```
hostname(config-cmap)# policy-map policy_map_name
hostname(config-pmap)#
```

where *policy\_map\_name* is the name of the policy map. The CLI enters the policy map configuration mode and the prompt changes accordingly.

- Step 5** Specify the class map, created in [Step 2](#), that identifies the traffic to be scanned. Use the **class** command to do so, as follows.

```
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

where *class\_map\_name* is the name of the class map you created in [Step 2](#). The CLI enters the policy map class configuration mode and the prompt changes accordingly.

- Step 6** If you want to enforce a per-client limit for simultaneous connections that the adaptive security appliance diverts to the CSC SSM, use the **set connection** command, as follows:

```
hostname(config-pmap-c)# set connection per-client-max n
```

where *n* is the maximum simultaneous connections the adaptive security appliance will allow per client. This prevents a single client from abusing the services of the CSC SSM or any server protected by the SSM, including prevention of attempts at DoS attacks on HTTP, FTP, POP3, or SMTP servers that the CSC SSM protects.

- Step 7** Assign the traffic identified by the class map as traffic to be sent to the CSC SSM. Use the **csc** command to do so, as follows.

```
hostname(config-pmap-c)# csc {fail-close | fail-open}
```

The **fail-close** and **fail-open** keywords control how the adaptive security appliance treats traffic when the CSC SSM is unavailable. For more information about the operating modes and failure behavior, see the [“About the CSC SSM”](#) section on page 22-5.

- Step 8** Use the **service-policy** command to apply the policy map globally or to a specific interface, as follows:

```
hostname(config-pmap-c)# service-policy policy_map_name [global | interface interface_ID]
hostname(config)#
```

where *policy\_map\_name* is the policy map you configured in [Step 4](#). If you want to apply the policy map to traffic on all the interfaces, use the **global** keyword. If you want to apply the policy map to traffic on a specific interface, use the **interface** *interface\_ID* option, where *interface\_ID* is the name assigned to the interface with the **nameif** command.

Only one global policy is allowed. You can override the global policy on an interface by applying a service policy to that interface. You can only apply one policy map to each interface.

The adaptive security appliance begins diverting traffic to the CSC SSM as specified.

Example 22-1 is based on the network shown in Figure 22-3. It creates two service policies. The first policy, `csc_out_policy`, is applied to the inside interface and uses the `csc_out` access list to ensure that all outbound requests for FTP and POP3 are scanned. The `csc_out` access list also ensures that HTTP connections from inside to networks on the outside interface are scanned but it includes a deny ACE to exclude HTTP connections from inside to servers on the DMZ network.

The second policy, `csc_in_policy`, is applied to the outside interface and uses the `csc_in` access list to ensure that requests for SMTP and HTTP originating on the outside interface and destined for the DMZ network are scanned by the CSC SSM. Scanning HTTP requests protects the web server from HTTP file uploads.

### Example 22-1 Service Policies for a Common CSC SSM Scanning Scenario

```
hostname(config)# access-list csc_out permit tcp 192.168.10.0 255.255.255.0 any eq 21
hostname(config)# access-list csc_out deny tcp 192.168.10.0 255.255.255.0 192.168.20.0 255.255.255.0 eq 80
hostname(config)# access-list csc_out permit tcp 192.168.10.0 255.255.255.0 any eq 80
hostname(config)# access-list csc_out permit tcp 192.168.10.0 255.255.255.0 any eq 110

hostname(config)# class-map csc_outbound_class
hostname(config-cmap)# match access-list csc_out

hostname(config)# policy-map csc_out_policy
hostname(config-pmap)# class csc_outbound_class
hostname(config-pmap-c)# csc fail-close

hostname(config)# service-policy csc_out_policy interface inside

hostname(config)# access-list csc_in permit tcp any 192.168.20.0 255.255.255.0 eq 25
hostname(config)# access-list csc_in permit tcp any 192.168.20.0 255.255.255.0 eq 80

hostname(config)# class-map csc_inbound_class
hostname(config-cmap)# match access-list csc_in

hostname(config)# policy-map csc_in_policy
hostname(config-pmap)# class csc_inbound_class
hostname(config-pmap-c)# csc fail-close

hostname(config)# service-policy csc_in_policy interface outside
```



#### Note

FTP inspection must be enabled for CSC SSM to scan files transferred by FTP. FTP inspection is enabled by default.

## Checking SSM Status

To check the status of an SSM, use the **show module** command.

The follow example output is from an adaptive security appliance with a CSC SSM installed. The Status field indicates the operational status of the SSM. An SSM operating normally has a status of “Up” in the output of the **show module** command. While the adaptive security appliance transfers an application image to the SSM, the Status field in the output reads “Recover”. For more information about possible statuses, see the entry for the **show module** command in the *Cisco Security Appliance Command Reference*.

```
hostname# show module 1
Mod Card Type Model Serial No.

```

## Transferring an Image onto an SSM

```

0 ASA 5520 Adaptive Security Appliance ASA5520 P3000000034
1 ASA 5500 Series Security Services Module-20 ASA-SSM-20 0

Mod MAC Address Range Hw Version Fw Version Sw Version

0 000b.fcf8.c30d to 000b.fcf8.c311 1.0 1.0(10)0 7.1(0)1
1 000b.fcf8.012c to 000b.fcf8.012c 1.0 1.0(10)0 Trend Micro InterScan Security Module Version 5.0

Mod SSM Application Name SSM Application Version

1 Trend Micro InterScan Security Version 5.0

Mod Status Data Plane Status Compatability

0 Up Sys Not Applicable
1 Up Up

```

The argument **1**, at the end of the command, is the slot number occupied by the SSM. If you do not know the slot number, you can omit it and see information about all modules, including the adaptive security appliance, which is considered to occupy slot 0 (zero).

Use the details keyword to view additional information for the SSM.

The follow example output is from an adaptive security appliance with a CSC SSM installed.

```

hostname# show module 1 details
Getting details from the Service Module, please wait...
ASA 5500 Series Security Services Module-20
Model: ASA-SSM-20
Hardware version: 1.0
Serial Number: 0
Firmware version: 1.0(10)0
Software version: Trend Micro InterScan Security Module Version 5.0
App. name: Trend Micro InterScan Security Module
App. version: Version 5.0
Data plane Status: Up
Status: Up
HTTP Service: Up
Mail Service: Up
FTP Service: Up
Activated: Yes
Mgmt IP addr: 10.23.62.92
Mgmt web port: 8443

```

## Transferring an Image onto an SSM

For an intelligent SSM, such as AIP SSM or CSC SSM, you can transfer application images from a TFTP server to the SSM. This process supports upgrade images and maintenance images.



### Note

If you are upgrading the application on the SSM, the SSM application may support backup of its configuration. If you do not back up the configuration of the SSM application, it is lost when you transfer an image onto the SSM. For more information about how your SSM supports backups, see the documentation for your SSM.

To transfer an image onto an intelligent SSM, perform the following steps:

**Step 1** Create or modify a recovery configuration for the SSM. To do so, perform the following steps:

- a. Determine if there is a recovery configuration for the SSM. To do so, use the **show module** command with the **recover** keyword, as follows.

```
hostname# show module slot recover
```

where *slot* is the slot number occupied by the SSM.

If the **recover** keyword is not valid, a recovery configuration does not exist. The **recover** keyword of the **show module** command is available only when a recovery configuration exists for the SSM.



**Note** When the adaptive security appliance operates in multiple context mode, the **configure** keyword is available only in the system context.

If there is a recovery configuration for the SSM, the adaptive security appliance displays it. Examine the recovery configuration closely to ensure that it is correct, especially the Image URL field. The following example shows a recovery configuration for an SSM in slot 1.

```
hostname# show module 1 recover
Module 1 recover parameters. . .
Boot Recovery Image: Yes
Image URL: tftp://10.21.18.1/ids-oldimg
Port IP Address: 10.1.2.10
Port Mask : 255.255.255.0
Gateway IP Address: 10.1.2.254
```

- b. If you need to create or modify the recovery configuration, use the **hw-module module recover** command with the **configure** keyword, as follows:

```
hostname# hw-module module slot recover configure
```

where *slot* is the slot number occupied by the SSM.

Complete the prompts as applicable. If you are modifying a configuration, you can keep the previously configured value by pressing **Enter**. The following example shows the prompts. For more information about them, see the entry for the **hw-module module recover** command in the *Cisco Security Appliance Command Line Configuration Guide*.

```
Image URL [tftp://0.0.0.0/]:
Port IP Address [0.0.0.0]:
VLAN ID [0]:
Gateway IP Address [0.0.0.0]:
```



**Note** Be sure the TFTP server you specify can transfer files up to 60 MB in size. Also, be sure the TFTP server can connect to the management port IP address that you specify for the SSM.

After you complete the prompts, the adaptive security appliance is ready to transfer to the SSM the image that it finds at the URL you specified.

**Step 2** Transfer the image from the TFTP server to the SSM and restart the SSM. To do so, use the **hw-module module recover** command with the **boot** keyword, as follows.

```
hostname# hw-module module slot recover boot
```

where *slot* is the slot number occupied by the SSM.

**Step 3** Check the progress of the image transfer and SSM restart process. To do so, use the **show module** command. For details, see the [“Checking SSM Status” section on page 22-13](#).

When the adaptive security appliance completes the image transfer and restart of the SSM, the SSM is running the newly transferred image.

**Note**

If your SSM supports configuration backups and you want to restore the configuration of the application running on the SSM, see the documentation for your SSM for details.

---



## CHAPTER 23

# Preventing Network Attacks

---

This chapter describes how to prevent network attacks by configuring TCP normalization, limiting TCP and UDP connections, and many other protection features.

This chapter includes the following sections:

- [Configuring TCP Normalization, page 23-1](#)
- [Configuring Connection Limits and Timeouts, page 23-6](#)
- [Preventing IP Spoofing, page 23-10](#)
- [Configuring the Fragment Size, page 23-11](#)
- [Blocking Unwanted Connections, page 23-11](#)
- [Configuring IP Audit for Basic IPS Support, page 23-12](#)

## Configuring TCP Normalization

The TCP normalization feature identifies abnormal packets that the security appliance can act on when they are detected; for example, the security appliance can allow, drop, or clear the packets. TCP normalization helps protect the security appliance from attacks. This section includes the following topics:

- [TCP Normalization Overview, page 23-1](#)
- [Enabling the TCP Normalizer, page 23-2](#)

## TCP Normalization Overview

The TCP normalizer includes non-configurable actions and configurable actions. Typically, non-configurable actions that drop or clear connections apply to packets that are always bad. Configurable actions (as detailed in [“Enabling the TCP Normalizer” section on page 23-2](#)) might need to be customized depending on your network needs.

See the following guidelines for TCP normalization:

- The normalizer does not protect from SYN floods. The security appliance includes SYN flood protection in other ways.
- The normalizer always sees the SYN packet as the first packet in a flow unless the security appliance is in loose mode due to failover.

## Enabling the TCP Normalizer

This feature uses Modular Policy Framework, so that implementing TCP normalization consists of identifying traffic, specifying the TCP normalization actions, and activating TCP normalization on an interface. See [Chapter 21, “Using Modular Policy Framework,”](#) for more information.

To configure TCP normalization, perform the following steps:

- Step 1** To specify the TCP normalization criteria that you want to look for, create a TCP map by entering the following command:

```
hostname(config)# tcp-map tcp-map-name
```

For each TCP map, you can customize one or more settings.

- Step 2** (Optional) Configure the TCP map criteria by entering one or more of the following commands (see [Table 23-1](#)). If you want to use the default settings for all criteria, you do not need to enter any commands for the TCP map. If you want to customize some settings, then the defaults are used for any commands you do not enter. The default configuration includes the following settings:

```
no check-retransmission
no checksum-verification
exceed-mss allow
queue-limit 0 timeout 4
reserved-bits allow
syn-data allow
synack-data drop
invalid-ack drop
seq-past-window drop
tcp-options range 6 7 clear
tcp-options range 9 255 clear
tcp-options selective-ack allow
tcp-options timestamp allow
tcp-options window-scale allow
ttl-evasion-protection
urgent-flag clear
window-variation allow-connection
```

**Table 23-1** *tcp-map Commands*

| Command                          | Notes                                                                                                                                                                                                                                                                                                     |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>check-retransmission</b>      | Prevents inconsistent TCP retransmissions.                                                                                                                                                                                                                                                                |
| <b>checksum-verification</b>     | Verifies the checksum.                                                                                                                                                                                                                                                                                    |
| <b>exceed-mss {allow   drop}</b> | Sets the action for packets whose data length exceeds the TCP maximum segment size.<br><br>(Default) The <b>allow</b> keyword allows packets whose data length exceeds the TCP maximum segment size.<br><br>The <b>drop</b> keyword drops packets whose data length exceeds the TCP maximum segment size. |



Table 23-1 *tcp-map Commands (continued)*

| Command                                                                | Notes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>invalid-ack</b> { <b>allow</b>   <b>drop</b> }                      | <p>Sets the action for packets with an invalid ACK. You might see invalid ACKs in the following instances:</p> <ul style="list-style-type: none"> <li>In the TCP connection SYN-ACK-received status, if the ACK number of a received TCP packet is not exactly same as the sequence number of the next TCP packet sending out, it is an invalid ACK.</li> <li>Whenever the ACK number of a received TCP packet is greater than the sequence number of the next TCP packet sending out, it is an invalid ACK.</li> </ul> <p>The <b>allow</b> keyword allows packets with an invalid ACK.</p> <p>(Default) The <b>drop</b> keyword drops packets with an invalid ACK.</p> <p><b>Note</b> TCP packets with an invalid ACK are automatically allowed for WAAS connections.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>queue-limit</b> <i>pkt_num</i><br>[ <b>timeout</b> <i>seconds</i> ] | <p>Sets the maximum number of out-of-order packets that can be buffered and put in order for a TCP connection, between 1 and 250 packets. The default is 0, which means this setting is disabled and the default system queue limit is used depending on the type of traffic:</p> <ul style="list-style-type: none"> <li>Connections for application inspection (the <b>inspect</b> command), IPS (the <b>ips</b> command), and TCP check-retransmission (the TCP map <b>check-retransmission</b> command) have a queue limit of 3 packets. If the security appliance receives a TCP packet with a different window size, then the queue limit is dynamically changed to match the advertised setting.</li> <li>For other TCP connections, out-of-order packets are passed through untouched.</li> </ul> <p>If you set the <b>queue-limit</b> command to be 1 or above, then the number of out-of-order packets allowed for all TCP traffic matches this setting. For application inspection, IPS, and TCP check-retransmission traffic, any advertised settings are ignored. For other TCP traffic, out-of-order packets are now buffered and put in order instead of passed through untouched.</p> <p>The <b>timeout</b> <i>seconds</i> argument sets the maximum amount of time that out-of-order packets can remain in the buffer, between 1 and 20 seconds; if they are not put in order and passed on within the timeout period, then they are dropped. The default is 4 seconds. You cannot change the timeout for any traffic if the <i>pkt_num</i> argument is set to 0; you need to set the limit to be 1 or above for the <b>timeout</b> keyword to take effect.</p> |

Table 23-1 *tcp-map Commands (continued)*

| Command                                                                                                                                                                                                                       | Notes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>reserved-bits</b> { <b>allow</b>   <b>clear</b>   <b>drop</b> }                                                                                                                                                            | <p>Sets the action for reserved bits in the TCP header.</p> <p>(Default) The <b>allow</b> keyword allows packets with the reserved bits in the TCP header.</p> <p>The <b>clear</b> keyword clears the reserved bits in the TCP header and allows the packet.</p> <p>The <b>drop</b> keyword drops the packet with the reserved bits in the TCP header.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>seq-past-window</b> { <b>allow</b>   <b>drop</b> }                                                                                                                                                                         | <p>Sets the action for packets that have past-window sequence numbers, namely the sequence number of a received TCP packet is greater than the right edge of the TCP receiving window.</p> <p>The <b>allow</b> keyword allows packets that have past-window sequence numbers. This action is only allowed if the <b>queue-limit</b> command is set to 0 (disabled).</p> <p>(Default) The <b>drop</b> keyword drops packets that have past-window sequence numbers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>synack-data</b> { <b>allow</b>   <b>drop</b> }                                                                                                                                                                             | <p>Sets the action for TCP SYNACK packets that contain data.</p> <p>The <b>allow</b> keyword allows TCP SYNACK packets that contain data.</p> <p>(Default) The <b>drop</b> keyword drops TCP SYNACK packets that contain data.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>syn-data</b> { <b>allow</b>   <b>drop</b> }                                                                                                                                                                                | <p>Sets the action for SYN packets with data.</p> <p>(Default) The <b>allow</b> keyword allows SYN packets with data.</p> <p>The <b>drop</b> keyword drops SYN packets with data.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>tcp-options</b> { <b>selective-ack</b>   <b>timestamp</b>   <b>window-scale</b> }<br>{ <b>allow</b>   <b>clear</b> }<br>Or<br><b>tcp-options range</b> <i>lower upper</i><br>{ <b>allow</b>   <b>clear</b>   <b>drop</b> } | <p>Sets the action for packets with TCP options, including the selective-ack, timestamp, or window-scale TCP options.</p> <p>(Default) The <b>allow</b> keyword allows packets with the specified option.</p> <p>(Default for <b>range</b>) The <b>clear</b> keyword clears the option and allows the packet.</p> <p>The <b>drop</b> keyword drops the packet with the specified option.</p> <p>The <b>selective-ack</b> keyword sets the action for the SACK option.</p> <p>The <b>timestamp</b> keyword sets the action for the timestamp option. Clearing the timestamp option disables PAWS and RTT.</p> <p>The <b>window-scale</b> keyword sets the action for the window scale mechanism option.</p> <p>The <b>range</b> keyword specifies a range of options. The <i>lower</i> argument sets the lower end of the range as 6, 7, or 9 through 255.</p> <p>The <i>upper</i> argument sets the upper end of the range as 6, 7, or 9 through 255.</p> |

**Table 23-1** *tcp-map Commands (continued)*

| Command                                | Notes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ttl-evasion-protection</b>          | <p>Disables the TTL evasion protection. Do not enter this command if you want to prevent attacks that attempt to evade security policy.</p> <p>For example, an attacker can send a packet that passes policy with a very short TTL. When the TTL goes to zero, a router between the security appliance and the endpoint drops the packet. It is at this point that the attacker can send a malicious packet with a long TTL that appears to the security appliance to be a retransmission and is passed. To the endpoint host, however, it is the first packet that has been received by the attacker. In this case, an attacker is able to succeed without security preventing the attack.</p> |
| <b>urgent-flag {allow   clear}</b>     | <p>Sets the action for packets with the URG flag. The URG flag is used to indicate that the packet contains information that is of higher priority than other data within the stream. The TCP RFC is vague about the exact interpretation of the URG flag, therefore end systems handle urgent offsets in different ways, which may make the end system vulnerable to attacks.</p> <p>The <b>allow</b> keyword allows packets with the URG flag.</p> <p>(Default) The <b>clear</b> keyword clears the URG flag and allows the packet.</p>                                                                                                                                                       |
| <b>window-variation {allow   drop}</b> | <p>Sets the action for a connection that has changed its window size unexpectedly. The window size mechanism allows TCP to advertise a large window and to subsequently advertise a much smaller window without having accepted too much data. From the TCP specification, “shrinking the window” is strongly discouraged. When this condition is detected, the connection can be dropped.</p> <p>(Default) The <b>allow</b> keyword allows connections with a window variation.</p> <p>The <b>drop</b> keyword drops connections with a window variation.</p>                                                                                                                                  |

- Step 3** To identify the traffic, add a class map using the **class-map** command. See the [“Creating a Layer 3/4 Class Map for Through Traffic”](#) section on page 21-5 for more information.

For example, you can match all traffic using the following commands:

```
hostname(config)# class-map TCPNORM
hostname(config-cmap)# match any
```

To match specific traffic, you can match an access list:

```
hostname(config)# access list TCPNORM extended permit ip any 10.1.1.1 255.255.255.255
hostname(config)# class-map TCP_norm_class
hostname(config-cmap)# match access-list TCPNORM
```

- Step 4** To add or edit a policy map that sets the actions to take with the class map traffic, enter the following commands:

```
hostname(config)# policy-map name
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

where the *class\_map\_name* is the class map from [Step 1](#).

For example:

```
hostname(config)# policy-map TCP_norm_policy
hostname(config-pmap)# class TCP_norm_class
hostname(config-pmap-c)#
```

**Step 5** Apply the TCP map to the class map by entering the following command.

```
hostname(config-pmap-c)# set connection advanced-options tcp-map-name
```

**Step 6** To activate the policy map on one or more interfaces, enter the following command:

```
hostname(config)# service-policy policymap_name {global | interface interface_name}
```

Where **global** applies the policy map to all interfaces, and **interface** applies the policy to one interface. Only one global policy is allowed. Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with inspections, and an interface policy with TCP normalization, then both inspections and TCP normalization are applied to the interface. However, if you have a global policy with inspections, and an interface policy with inspections, then only the interface policy inspections are applied to that interface.

For example, to allow urgent flag and urgent offset packets for all traffic sent to the range of TCP ports between the well known FTP data port and the Telnet port, enter the following commands:

```
hostname(config)# tcp-map tmap
hostname(config-tcp-map)# urgent-flag allow
hostname(config-tcp-map)# class-map urg-class
hostname(config-cmap)# match port tcp range ftp-data telnet
hostname(config-cmap)# policy-map pmap
hostname(config-pmap)# class urg-class
hostname(config-pmap-c)# set connection advanced-options tmap
hostname(config-pmap-c)# service-policy pmap global
```

## Configuring Connection Limits and Timeouts

This section describes how to set maximum TCP and UDP connections, maximum embryonic connections, maximum per-client connections, connection timeouts, dead connection detection, and how to disable TCP sequence randomization. You can set limits for connections that go through the security appliance, or for management connections to the security appliance. This section includes the following topics:

- [Connection Limit Overview, page 23-7](#)
- [Enabling Connection Limits and Timeouts, page 23-8](#)



### Note

You can also configure maximum connections, maximum embryonic connections, and TCP sequence randomization in the NAT configuration. If you configure these settings for the same traffic using both methods, then the security appliance uses the lower limit. For TCP sequence randomization, if it is disabled using either method, then the security appliance disables TCP sequence randomization.

## Connection Limit Overview

This section describes why you might want to limit connections, and includes the following topics:

- [TCP Intercept Overview, page 23-7](#)
- [Disabling TCP Intercept for Management Packets for Clientless SSL Compatibility, page 23-7](#)
- [Dead Connection Detection \(DCD\) Overview, page 23-7](#)
- [TCP Sequence Randomization Overview, page 23-8](#)

### TCP Intercept Overview

Limiting the number of embryonic connections protects you from a DoS attack. The security appliance uses the per-client limits and the embryonic connection limit to trigger TCP Intercept, which protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. TCP Intercept uses the SYN cookies algorithm to prevent TCP SYN-flooding attacks. A SYN-flooding attack consists of a series of SYN packets usually originating from spoofed IP addresses. The constant flood of SYN packets keeps the server SYN queue full, which prevents it from servicing connection requests. When the embryonic connection threshold of a connection is crossed, the security appliance acts as a proxy for the server and generates a SYN-ACK response to the client SYN request. When the security appliance receives an ACK back from the client, it can then authenticate the client and allow the connection to the server.

### Disabling TCP Intercept for Management Packets for Clientless SSL Compatibility

By default, TCP management connections have TCP Intercept always enabled. When TCP Intercept is enabled, it intercepts the 3-way TCP connection establishment handshake packets and thus deprives the security appliance from processing the packets for clientless SSL. Clientless SSL requires the ability to process the 3-way handshake packets to provide selective ACK and other TCP options for clientless SSL connections. To disable TCP Intercept for management traffic, you can set the embryonic connection limit; only after the embryonic connection limit is reached is TCP Intercept enabled.

### Dead Connection Detection (DCD) Overview

DCD detects a dead connection and allows it to expire, without expiring connections that can still handle traffic. You configure DCD when you want idle, but valid connections to persist.

When you enable DCD, idle timeout behavior changes. With idle timeout, DCD probes are sent to each of the two end-hosts to determine the validity of the connection. If an end-host fails to respond after probes are sent at the configured intervals, the connection is freed, and reset values, if configured, are sent to each of the end-hosts. If both end-hosts respond that the connection is valid, the activity timeout is updated to the current time and the idle timeout is rescheduled accordingly.

Enabling DCD changes the behavior of idle-timeout handling in the TCP normalizer. DCD probing resets the idle timeout on the connections seen in the **show conn** command. To determine when a connection that has exceeded the configured timeout value in the timeout command but is kept alive due to DCD probing, the **show service-policy** command includes counters to show the amount of activity from DCD.

## TCP Sequence Randomization Overview

Each TCP connection has two ISNs: one generated by the client and one generated by the server. The security appliance randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

TCP initial sequence number randomization can be disabled if required. For example:

- If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
- If you use eBGP multi-hop through the security appliance, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
- You use a WAAS device that requires the security appliance not to randomize the sequence numbers of connections.

## Enabling Connection Limits and Timeouts

To set connection limits and timeouts, perform the following steps:

- Step 1** To identify the traffic, add a class map using the **class-map** command. See the [“Creating a Layer 3/4 Class Map for Through Traffic” section on page 21-5](#) for more information.

For example, you can match all traffic using the following commands:

```
hostname(config)# class-map CONNS
hostname(config-cmap)# match any
```

To match specific traffic, you can match an access list:

```
hostname(config)# access list CONNS extended permit ip any 10.1.1.1 255.255.255.255
hostname(config)# class-map CONNS
hostname(config-cmap)# match access-list CONNS
```

- Step 2** To add or edit a policy map that sets the actions to take with the class map traffic, enter the following commands:

```
hostname(config)# policy-map name
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

where the *class\_map\_name* is the class map from [Step 1](#).

For example:

```
hostname(config)# policy-map CONNS
hostname(config-pmap)# class CONNS
hostname(config-pmap-c)#
```

- Step 3** To set maximum connection limits or whether TCP sequence randomization is enabled, enter the following command:

```
hostname(config-pmap-c)# set connection {[conn-max n] [embryonic-conn-max n]
[per-client-embryonic-max n] [per-client-max n] [random-sequence-number {enable |
disable}]}
```

where the **conn-max** *n* argument sets the maximum number of simultaneous TCP and/or UDP connections that are allowed, between 0 and 65535. The default is 0, which allows unlimited connections.

The **embryonic-conn-max** *n* argument sets the maximum number of simultaneous embryonic connections allowed, between 0 and 65535. The default is 0, which allows unlimited connections.

The **per-client-embryonic-max** *n* argument sets the maximum number of simultaneous embryonic connections allowed per client, between 0 and 65535. The default is 0, which allows unlimited connections.

The **per-client-max** *n* argument sets the maximum number of simultaneous connections allowed per client, between 0 and 65535. The default is 0, which allows unlimited connections.

The **random-sequence-number** {**enable** | **disable**} keyword enables or disables TCP sequence number randomization. See the “TCP Sequence Randomization Overview” section on page 23-8 section for more information.

You can enter this command all on one line (in any order), or you can enter each attribute as a separate command. The security appliance combines the command into one line in the running configuration.

**Step 4** To set connection timeouts, enter the following command:

```
hostname(config-pmap-c)# set connection timeout {[embryonic hh:mm:ss] {tcp hh:mm:ss
[reset]] [half-closed hh:mm:ss] [dcd hh:mm:ss [max_retries]]}
```

where the **embryonic** *hh:mm:ss* keyword sets the timeout period until a TCP embryonic (half-open) connection is closed, between 0:0:5 and 1193:00:00. The default is 0:0:30. You can also set this value to 0, which means the connection never times out.

The **tcp** *hh:mm:ss* keyword sets the idle timeout between 0:5:0 and 1193:00:00. The default is 1:0:0. You can also set this value to 0, which means the connection never times out. The **reset** keyword sends a reset to TCP endpoints when the connection times out. The security appliance sends the reset packet only in response to a host sending another packet for the timed-out flow (on the same source and destination port). The host then removes the connection from its connection table after receiving the reset packet. The host application can then attempt to establish a new connection using a SYN packet.

The **half-closed** *hh:mm:ss* keyword sets the idle timeout between 0:5:0 and 1193:00:00. The default is 0:10:0. Half-closed connections are not affected by DCD. Also, the security appliance does not send a reset when taking down half-closed connections.

The **dcd** keyword enables DCD. DCD detects a dead connection and allows it to expire, without expiring connections that can still handle traffic. You configure DCD when you want idle, but valid connections to persist. After a TCP connection times out, the security appliance sends DCD probes to the end hosts to determine the validity of the connection. If one of the end hosts fails to respond after the maximum retries are exhausted, the security appliance frees the connection. If both end hosts respond that the connection is valid, the security appliance updates the activity timeout to the current time and reschedules the idle timeout accordingly. The *retry-interval* sets the time duration in *hh:mm:ss* format to wait after each unresponsive DCD probe before sending another probe, between 0:0:1 and 24:0:0. The default is 0:0:15. The *max-retries* sets the number of consecutive failed retries for DCD before declaring the connection as dead. The minimum value is 1 and the maximum value is 255. The default is 5.

You can enter this command all on one line (in any order), or you can enter each attribute as a separate command. The command is combined onto one line in the running configuration.

**Step 5** To activate the policy map on one or more interfaces, enter the following command:

```
hostname(config)# service-policy polycmap_name {global | interface interface_name}
```

Where **global** applies the policy map to all interfaces, and **interface** applies the policy to one interface. Only one global policy is allowed. Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with inspections, and an interface policy with TCP normalization, then both inspections and TCP normalization are applied to the interface. However, if you have a global policy with inspections, and an interface policy with inspections, then only the interface policy inspections are applied to that interface.

The following example sets the connection limits and timeouts for all traffic:

```
hostname(config)# class-map CONNS
hostname(config-cmap)# match any
hostname(config-cmap)# policy-map CONNS
hostname(config-pmap)# class CONNS
hostname(config-pmap-c)# set connection conn-max 1000 embryonic-conn-max 3000
hostname(config-pmap-c)# set connection timeout tcp 2:0:0 embryonic 0:40:0 half-closed
0:20:0 dcd
hostname(config-pmap-c)# service-policy CONNS interface outside
```

You can enter **set connection** commands with multiple parameters or you can enter each parameter as a separate command. The security appliance combines the commands into one line in the running configuration. For example, if you entered the following two commands in class configuration mode:

```
hostname(config-pmap-c)# set connection conn-max 600
hostname(config-pmap-c)# set connection embryonic-conn-max 50
```

the output of the **show running-config policy-map** command would display the result of the two commands in a single, combined command:

```
set connection conn-max 600 embryonic-conn-max 50
```

## Preventing IP Spoofing

This section lets you enable Unicast Reverse Path Forwarding on an interface. Unicast RPF guards against IP spoofing (a packet uses an incorrect source IP address to obscure its true source) by ensuring that all packets have a source IP address that matches the correct source interface according to the routing table.

Normally, the security appliance only looks at the destination address when determining where to forward the packet. Unicast RPF instructs the security appliance to also look at the source address; this is why it is called Reverse Path Forwarding. For any traffic that you want to allow through the security appliance, the security appliance routing table must include a route back to the source address. See RFC 2267 for more information.

For outside traffic, for example, the security appliance can use the default route to satisfy the Unicast RPF protection. If traffic enters from an outside interface, and the source address is not known to the routing table, the security appliance uses the default route to correctly identify the outside interface as the source interface.

If traffic enters the outside interface from an address that is known to the routing table, but is associated with the inside interface, then the security appliance drops the packet. Similarly, if traffic enters the inside interface from an unknown source address, the security appliance drops the packet because the matching route (the default route) indicates the outside interface.

Unicast RPF is implemented as follows:

- ICMP packets have no session, so each packet is checked.



- UDP and TCP have sessions, so the initial packet requires a reverse route lookup. Subsequent packets arriving during the session are checked using an existing state maintained as part of the session. Non-initial packets are checked to ensure they arrived on the same interface used by the initial packet.

To enable Unicast RPF, enter the following command:

```
hostname(config)# ip verify reverse-path interface interface_name
```

## Configuring the Fragment Size

By default, the security appliance allows up to 24 fragments per IP packet, and up to 200 fragments awaiting reassembly. You might need to let fragments on your network if you have an application that routinely fragments packets, such as NFS over UDP. However, if you do not have an application that fragments traffic, we recommend that you do not allow fragments through the security appliance. Fragmented packets are often used as DoS attacks. To set disallow fragments, enter the following command:

```
hostname(config)# fragment chain 1 [interface_name]
```

Enter an interface name if you want to prevent fragmentation on a specific interface. By default, this command applies to all interfaces.

## Blocking Unwanted Connections

If you know that a host is attempting to attack your network (for example, system log messages show an attack), then you can block (or shun) connections based on the source IP address and other identifying parameters. No new connections can be made until you remove the shun.



### Note

If you have an IPS that monitors traffic, such as an AIP SSM, then the IPS can shun connections automatically.

To shun a connection manually, perform the following steps:

- Step 1** If necessary, view information about the connection by entering the following command:

```
hostname# show conn
```

The security appliance shows information about each connection, such as the following:

```
TCP out 64.101.68.161:4300 in 10.86.194.60:23 idle 0:00:00 bytes 1297 flags UIO
```

- Step 2** To shun connections from the source IP address, enter the following command:

```
hostname(config)# shun src_ip [dst_ip src_port dest_port [protocol]] [vlan vlan_id]
```

If you enter only the source IP address, then all future connections are shunned; existing connections remain active.

To drop an existing connection, as well as blocking future connections from the source IP address, enter the destination IP address, source and destination ports, and the protocol. By default, the protocol is 0 for IP.

For multiple context mode, you can enter this command in the admin context, and by specifying a VLAN ID that is assigned to an interface in other contexts, you can shun the connection in other contexts.

**Step 3** To remove the shun, enter the following command:

```
hostname(config)# no shun src_ip [vlan vlan_id]
```

---

## Configuring IP Audit for Basic IPS Support

The IP audit feature provides basic IPS support for a security appliance that does not have an AIP SSM. It supports a basic list of signatures, and you can configure the security appliance to perform one or more actions on traffic that matches a signature.

To enable IP audit, perform the following steps:

**Step 1** To define an IP audit policy for informational signatures, enter the following command:

```
hostname(config)# ip audit name name info [action [alarm] [drop] [reset]]
```

Where **alarm** generates a system message showing that a packet matched a signature, **drop** drops the packet, and **reset** drops the packet and closes the connection. If you do not define an action, then the default action is to generate an alarm.

**Step 2** To define an IP audit policy for attack signatures, enter the following command:

```
hostname(config)# ip audit name name attack [action [alarm] [drop] [reset]]
```

Where **alarm** generates a system message showing that a packet matched a signature, **drop** drops the packet, and **reset** drops the packet and closes the connection. If you do not define an action, then the default action is to generate an alarm.

**Step 3** To assign the policy to an interface, enter the following command:

```
ip audit interface interface_name policy_name
```

**Step 4** To disable signatures, or for more information about signatures, see the **ip audit signature** command in the *Cisco Security Appliance Command Line Configuration Guide*.

---



# CHAPTER 24

## Configuring QoS

---

Have you ever participated in a long-distance phone call that involved a satellite connection? The conversation might be interrupted with brief, but perceptible, gaps at odd intervals. Those gaps are the time, called the latency, between the arrival of packets being transmitted over the network. Some network traffic, such as voice and video, cannot tolerate long latency times. Quality of Service (QoS) is a feature that lets you give priority to critical traffic, prevent bandwidth hogging, and manage network bottlenecks to prevent packet drops.

This chapter describes how to apply QoS policies, and includes the following sections:

- [QoS Overview, page 24-1](#)
- [Creating the Standard Priority Queue for an Interface, page 24-5](#)
- [Identifying Traffic for QoS Using Class Maps, page 24-8](#)
- [Creating a Policy for Standard Priority Queueing and/or Policing, page 24-9](#)
- [Creating a Policy for Traffic Shaping and Hierarchical Priority Queueing, page 24-11](#)
- [Viewing QoS Statistics, page 24-13](#)

## QoS Overview

You should consider that in an ever-changing network environment, QoS is not a one-time deployment, but an ongoing, essential part of network design.



### Note

---

QoS is only available in single context mode.

---

This section describes the QoS features supported by the security appliance, and includes the following topics:

- [Supported QoS Features, page 24-2](#)
- [What is a Token Bucket?, page 24-2](#)
- [Policing Overview, page 24-3](#)
- [Priority Queueing Overview, page 24-3](#)
- [Traffic Shaping Overview, page 24-4](#)
- [DSCP and DiffServ Preservation, page 24-5](#)

## Supported QoS Features

The security appliance supports the following QoS features:

- **Policing**—To prevent individual flows from hogging the network bandwidth, you can limit the maximum bandwidth used per flow. See the [“Policing Overview” section on page 24-3](#) for more information.
- **Priority queuing**—For critical traffic that cannot tolerate latency, such as Voice over IP (VoIP), you can identify traffic for Low Latency Queuing (LLQ) so that it is always transmitted ahead of other traffic. See the [“Priority Queueing Overview” section on page 24-3](#) for more information.
- **Traffic shaping**—If you have a device that transmits packets at a high speed, such as a security appliance with Fast Ethernet, and it is connected to a low speed device such as a cable modem, then the cable modem is a bottleneck at which packets are frequently dropped. To manage networks with differing line speeds, you can configure the security appliance to transmit packets at a fixed slower rate. See the [“Traffic Shaping Overview” section on page 24-4](#) for more information.

## What is a Token Bucket?

A token bucket is used to manage a device that regulates the data in a flow. For example, the regulator might be a traffic policer or a traffic shaper. A token bucket itself has no discard or priority policy. Rather, a token bucket discards tokens and leaves to the flow the problem of managing its transmission queue if the flow overdrives the regulator.

A token bucket is a formal definition of a rate of transfer. It has three components: a burst size, an average rate, and a time interval. Although the average rate is generally represented as bits per second, any two values may be derived from the third by the relation shown as follows:

average rate = burst size / time interval

Here are some definitions of these terms:

- **Average rate**—Also called the committed information rate (CIR), it specifies how much data can be sent or forwarded per unit time on average.
- **Burst size**—Also called the Committed Burst (Bc) size, it specifies in bits or bytes per burst how much traffic can be sent within a given unit of time to not create scheduling concerns. (For traffic shaping, it specifies bits per burst; for policing, it specifies bytes per burst.)
- **Time interval**—Also called the measurement interval, it specifies the time quantum in seconds per burst.

In the token bucket metaphor, tokens are put into the bucket at a certain rate. The bucket itself has a specified capacity. If the bucket fills to capacity, newly arriving tokens are discarded. Each token is permission for the source to send a certain number of bits into the network. To send a packet, the regulator must remove from the bucket a number of tokens equal in representation to the packet size.

If not enough tokens are in the bucket to send a packet, the packet either waits until the bucket has enough tokens (in the case of traffic shaping) or the packet is discarded or marked down (in the case of policing). If the bucket is already full of tokens, incoming tokens overflow and are not available to future packets. Thus, at any time, the largest burst a source can send into the network is roughly proportional to the size of the bucket.

Note that the token bucket mechanism used for traffic shaping has both a token bucket and a data buffer, or queue; if it did not have a data buffer, it would be a policer. For traffic shaping, packets that arrive that cannot be sent immediately are delayed in the data buffer.

For traffic shaping, a token bucket permits burstiness but bounds it. It guarantees that the burstiness is bounded so that the flow will never send faster than the token bucket capacity, divided by the time interval, plus the established rate at which tokens are placed in the token bucket. See the following formula:

$$(\text{token bucket capacity in bits} / \text{time interval in seconds}) + \text{established rate in bps} = \text{maximum flow speed in bps}$$

This method of bounding burstiness also guarantees that the long-term transmission rate will not exceed the established rate at which tokens are placed in the bucket.

## Policing Overview

Policing is a way of ensuring that no traffic exceeds the maximum rate (in bits/second) that you configure, thus ensuring that no one traffic flow or class can take over the entire resource. When traffic exceeds the maximum rate, the security appliance drops the excess traffic. Policing also sets the largest single burst of traffic allowed.

## Priority Queueing Overview

LLQ priority queueing lets you prioritize certain traffic flows (such as latency-sensitive traffic like voice and video) ahead of other traffic.

The security appliance supports two types of priority queueing:

- **Standard priority queueing**—Standard priority queueing uses an LLQ priority queue on an interface (see the [“Creating the Standard Priority Queue for an Interface”](#) section on page 24-5), while all other traffic goes into the “best effort” queue. Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped. This is called *tail drop*. To avoid having the queue fill up, you can increase the queue buffer size. You can also fine-tune the maximum number of packets allowed into the transmit queue. These options let you control the latency and robustness of the priority queueing. Packets in the LLQ queue are always transmitted before packets in the best effort queue.
- **Hierarchical priority queueing**—Hierarchical priority queueing is used on interfaces on which you enable a traffic shaping queue. A subset of the shaped traffic can be prioritized. The standard priority queue is not used. See the following guidelines about hierarchical priority queueing:
  - Priority packets are always queued at the head of the shape queue so they are always transmitted ahead of other non-priority queued packets.
  - Priority packets are never dropped from the shape queue unless the sustained rate of priority traffic exceeds the shape rate.
  - For IPSec-encrypted packets, you can only match traffic based on the DSCP or precedence setting.
  - IPSec-over-TCP is not supported for priority traffic classification.

## Traffic Shaping Overview

Traffic shaping is used to match device and link speeds, thereby controlling packet loss, variable delay, and link saturation, which can cause jitter and delay.

- Traffic shaping must be applied to all outgoing traffic on a physical interface or in the case of the ASA 5505, on a VLAN. You cannot configure traffic shaping for specific types of traffic.
- Traffic shaping is implemented when packets are ready to be transmitted on an interface, so the rate calculation is performed based on the actual size of a packet to be transmitted, including all the possible overhead such as the IPSec header and L2 header.
- The shaped traffic includes both through-the-box and from-the-box traffic.
- The shape rate calculation is based on the standard token bucket algorithm. The token bucket size is twice the Burst Size value. See the [“What is a Token Bucket?”](#) section on page 24-2.
- When bursty traffic exceeds the specified shape rate, packets are queued and transmitted later. Following are some characteristics regarding the shape queue (for information about hierarchical priority queueing, see the [“Priority Queueing Overview”](#) section on page 24-3):
  - The queue size is calculated based on the shape rate. The queue can hold the equivalent of 200-milliseconds worth of shape rate traffic, assuming a 1500-byte packet. The minimum queue size is 64.
  - When the queue limit is reached, packets are tail-dropped.
  - Certain critical keep-alive packets such as OSPF Hello packets are never dropped.
  - The time interval is derived by  $time\_interval = burst\_size / average\_rate$ . The larger the time interval is, the burstier the shaped traffic might be, and the longer the link might be idle. The effect can be best understood using the following exaggerated example:

Average Rate = 1000000

Burst Size = 1000000

In the above example, the time interval is 1 second, which means, 1 Mbps of traffic can be bursted out within the first 10 milliseconds of the 1-second interval on a 100 Mbps FE link and leave the remaining 990 milliseconds idle without being able to send any packets until the next time interval. So if there is delay-sensitive traffic such as voice traffic, the Burst Size should be reduced compared to the average rate so the time interval is reduced.

## How QoS Features Interact

You can configure each of the QoS features alone if desired for the security appliance. Often, though, you configure multiple QoS features on the security appliance so you can prioritize some traffic, for example, and prevent other traffic from causing bandwidth problems.

See the following supported feature combinations per interface:

- Standard priority queuing (for specific traffic) + Policing (for the rest of the traffic).  
You cannot configure priority queueing and policing for the same set of traffic.
- Traffic shaping (for all traffic on an interface) + Hierarchical priority queueing (for a subset of traffic).

You cannot configure traffic shaping and standard priority queueing for the same interface; only hierarchical priority queueing is allowed. For example, if you configure standard priority queueing for the global policy, and then configure traffic shaping for a specific interface, the feature you configured last is rejected because the global policy overlaps the interface policy.

Typically, if you enable traffic shaping, you do not also enable policing for the same traffic, although the security appliance does not restrict you from configuring this.

## DSCP and DiffServ Preservation

- DSCP markings are preserved on all traffic passing through the security appliance.
- The security appliance does not locally mark/re-mark any classified traffic, but it honors the Expedited Forwarding (EF) DSCP bits of every packet to determine if it requires “priority” handling and will direct those packets to the LLQ.
- DiffServ marking is preserved on packets when they traverse the service provider backbone so that QoS can be applied in transit (QoS tunnel pre-classification).

## Creating the Standard Priority Queue for an Interface

If you enable standard priority queueing for traffic on a physical interface, then you need to also create the priority queue on each interface. Each physical interface uses two queues: one for priority traffic, and the other for all other traffic. For the other traffic, you can optionally configure policing.



### Note

The standard priority queue is not required for hierarchical priority queueing with traffic shaping; see the [“Priority Queueing Overview”](#) section on page 24-3 for more information.

This section includes the following topics:

- [Determining the Queue and TX Ring Limits](#), page 24-6
- [Configuring the Priority Queue](#), page 24-7

## Determining the Queue and TX Ring Limits

To determine the priority queue and TX ring limits, use the worksheets below.

Table 24-1 shows how to calculate the priority queue size. Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped (called *tail drop*). To avoid having the queue fill up, you can adjust the queue buffer size according to the “Configuring the Priority Queue” section on page 24-7.

**Table 24-1 Queue Limit Worksheet**

|               |                                                |   |                                          |   |                            |
|---------------|------------------------------------------------|---|------------------------------------------|---|----------------------------|
| <b>Step 1</b> | _____ Mbps                                     | x | 125                                      | = | _____                      |
|               | Outbound bandwidth (Mbps or Kbps) <sup>1</sup> |   |                                          |   | # of bytes/ms              |
|               | _____ Kbps                                     | x | .125                                     | = | _____                      |
|               |                                                |   |                                          |   | # of bytes/ms              |
| <b>Step 2</b> | _____                                          | ÷ | _____                                    | x | _____                      |
|               | # of bytes/ms from Step 1                      |   | Average packet size (bytes) <sup>2</sup> |   | Delay (ms) <sup>3</sup>    |
|               |                                                |   |                                          | = | _____                      |
|               |                                                |   |                                          |   | Queue limit (# of packets) |

1. For example, DSL might have an uplink speed of 768 Kbps. Check with your provider.
2. Determine this value from a codec or sampling size. For example, for VoIP over VPN, you might use 160 bytes. We recommend 256 bytes if you do not know what size to use.
3. The delay depends on your application. For example, the recommended maximum delay for VoIP is 200 ms. We recommend 500 ms if you do not know what delay to use.

Table 24-2 shows how to calculate the TX ring limit. This limit determines the maximum number of packets allowed into the Ethernet transmit driver before the driver pushes back to the queues on the interface to let them buffer packets until the congestion clears. This setting guarantees that the hardware-based transmit ring imposes a limited amount of extra latency for a high-priority packet.

**Table 24-2 TX Ring Limit Worksheet**

|               |                                                |   |                                          |   |                              |
|---------------|------------------------------------------------|---|------------------------------------------|---|------------------------------|
| <b>Step 1</b> | _____ Mbps                                     | x | 125                                      | = | _____                        |
|               | Outbound bandwidth (Mbps or Kbps) <sup>1</sup> |   |                                          |   | # of bytes/ms                |
|               | _____ Kbps                                     | x | 0.125                                    | = | _____                        |
|               |                                                |   |                                          |   | # of bytes/ms                |
| <b>Step 2</b> | _____                                          | ÷ | _____                                    | x | _____                        |
|               | # of bytes/ms from Step 1                      |   | Maximum packet size (bytes) <sup>2</sup> |   | Delay (ms) <sup>3</sup>      |
|               |                                                |   |                                          | = | _____                        |
|               |                                                |   |                                          |   | TX ring limit (# of packets) |

1. For example, DSL might have an uplink speed of 768 Kbps. Check with your provider.
2. Typically, the maximum size is 1538 bytes, or 1542 bytes for tagged Ethernet.
3. The delay depends on your application. For example, to control jitter for VoIP, you should use 20 ms.



## Configuring the Priority Queue

To create the priority queue, perform the following steps.

- Step 1** To create the priority queue, enter the following command:

```
hostname(config)# priority-queue interface_name
```

Where the *interface\_name* argument specifies the physical interface name on which you want to enable the priority queue, or for the ASA 5505, the VLAN interface name.

- Step 2** (Optional) To change the size of the priority queues, enter the following command:

```
hostname(config-priority-queue)# queue-limit number_of_packets
```

The default queue limit is 1024 packets. Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped (called *tail drop*). To avoid having the queue fill up, you can use the **queue-limit** command to increase the queue buffer size.

See the [“Determining the Queue and TX Ring Limits” section on page 24-6](#) to determine the *number\_of\_packets* value.

The upper limit of the range of values for the **queue-limit** command is determined dynamically at run time. To view this limit, enter **queue-limit ?** on the command line. The key determinants are the memory needed to support the queues and the memory available on the device.

The **queue-limit** that you specify affects both the higher priority low-latency queue and the best effort queue.

- Step 3** (Optional) To specify the depth of the priority queues, enter the following command:

```
hostname(config-priority-queue)# tx-ring-limit number_of_packets
```

The default tx-ring-limit is 128 packets. This command sets the maximum number of low-latency or normal priority packets allowed into the Ethernet transmit driver before the driver pushes back to the queues on the interface to let them buffer packets until the congestion clears. This setting guarantees that the hardware-based transmit ring imposes a limited amount of extra latency for a high-priority packet.

See the [“Determining the Queue and TX Ring Limits” section on page 24-6](#) to determine the *number\_of\_packets* value.

The upper limit of the range of values for the **tx-ring-limit** command is determined dynamically at run time. To view this limit, enter **tx-ring-limit ?** on the command line. The key determinants are the memory needed to support the queues and the memory available on the device.

The **tx-ring-limit** that you specify affects both the higher priority low-latency queue and the best-effort queue.

The following example establishes a priority queue on interface “outside” (the GigabitEthernet0/1 interface), with the default queue-limit and tx-ring-limit.

```
hostname(config)# priority-queue outside
```

The following example establishes a priority queue on the interface “outside” (the GigabitEthernet0/1 interface), sets the queue-limit to 260 packets, and sets the tx-ring-limit to 3:

```
hostname(config)# priority-queue outside
hostname(config-priority-queue)# queue-limit 260
hostname(config-priority-queue)# tx-ring-limit 3
```

# Identifying Traffic for QoS Using Class Maps

QoS is part of the Modular Policy Framework. See the [Chapter 21, “Using Modular Policy Framework,”](#) for more information. In Modular Policy Framework, you identify the traffic on which you want to enable QoS in a class map. This section includes the following topics:

- [Creating a QoS Class Map, page 24-8](#)
- [QoS Class Map Examples, page 24-8](#)

## Creating a QoS Class Map

For priority traffic, identify only latency-sensitive traffic. For policing traffic, you can choose to police all other traffic, or you can limit the traffic to certain types. For traffic shaping, all traffic on an interface must be shaped.

To create the class maps for QoS traffic, see the **class-map** command in the [“Identifying Traffic \(Layer 3/4 Class Map\)”](#) section on page 21-4.

You can match traffic based on many characteristics, including access lists, tunnel groups, DSCP, precedence, and more. See the following guidelines for configuring class maps for QoS:

- For traffic shaping, you can only use the **class-default** class map, which is automatically created by the security appliance, and which matches all traffic.
- You cannot use the **class-default** class map for priority traffic.
- For hierarchical priority queueing, for IPSec-encrypted packets, you can only match traffic based on the DSCP or precedence setting.
- For hierarchical priority queueing, IPSec-over-TCP traffic is not supported.

## QoS Class Map Examples

For example, in the following sequence, the **class-map** command classifies all non-tunneled TCP traffic, using an access list named `tcp_traffic`:

```
hostname(config)# access-list tcp_traffic permit tcp any any
hostname(config)# class-map tcp_traffic
hostname(config-cmap)# match access-list tcp_traffic
```

In the following example, other, more specific match criteria are used for classifying traffic for specific, security-related tunnel groups. These specific match criteria stipulate that a match on tunnel-group (in this case, the previously-defined Tunnel-Group-1) is required as the first match characteristic to classify traffic for a specific tunnel, and it allows for an additional match line to classify the traffic (IP differential services code point, expedited forwarding).

```
hostname(config)# class-map TGl-voice
hostname(config-cmap)# match tunnel-group tunnel-grp1
hostname(config-cmap)# match dscp ef
```

In the following example, the **class-map** command classifies both tunneled and non-tunneled traffic according to the traffic type:

```
hostname(config)# access-list tunneled extended permit ip 10.10.34.0 255.255.255.0
20.20.10.0 255.255.255.0
hostname(config)# access-list non-tunneled extended permit tcp any any
hostname(config)# tunnel-group tunnel-grp1 type IPSec_L2L
```

```

hostname(config)# class-map browse
hostname(config-cmap)# description "This class-map matches all non-tunneled tcp traffic."
hostname(config-cmap)# match access-list non-tunneled

hostname(config-cmap)# class-map TG1-voice
hostname(config-cmap)# description "This class-map matches all dscp ef traffic for
tunnel-grp 1."
hostname(config-cmap)# match dscp ef
hostname(config-cmap)# match tunnel-group tunnel-grp1

hostname(config-cmap)# class-map TG1-BestEffort
hostname(config-cmap)# description "This class-map matches all best-effort traffic for
tunnel-grp1."
hostname(config-cmap)# match tunnel-group tunnel-grp1
hostname(config-cmap)# match flow ip destination-address

```

The following example shows a way of policing a flow within a tunnel, provided the classed traffic is not specified as a tunnel, but does go *through* the tunnel. In this example, 192.168.10.10 is the address of the host machine on the private side of the remote tunnel, and the access list is named “host-over-l2l”. By creating a class-map (named “host-specific”), you can then police the “host-specific” class before the LAN-to-LAN connection polices the tunnel. In this example, the “host-specific” traffic is rate-limited before the tunnel, then the tunnel is rate-limited:

```

hostname(config)# access-list host-over-l2l extended permit ip any host 192.168.10.10
hostname(config)# class-map host-specific
hostname(config-cmap)# match access-list host-over-l2l

```

The following example builds on the configuration developed in the previous section. As in the previous example, there are two named class-maps: tcp\_traffic and TG1-voice.

```

hostname(config)# class-map TG1-best-effort
hostname(config-cmap)# match tunnel-group Tunnel-Group-1
hostname(config-cmap)# match flow ip destination-address

```

Adding a third class map provides a basis for defining a tunneled and non-tunneled QoS policy, as follows, which creates a simple QoS policy for tunneled and non-tunneled traffic, assigning packets of the class TG1-voice to the low latency queue and setting rate limits on the tcp\_traffic and TG1-best-effort traffic flows.

## Creating a Policy for Standard Priority Queueing and/or Policing

After you identify the traffic in [“Identifying Traffic for QoS Using Class Maps” section on page 24-8](#), you can create a policy map for an interface or globally for all interfaces that assigns QoS actions (and other feature actions) to the traffic in the class map. (See the [Chapter 21, “Using Modular Policy Framework,”](#) for information about other features. This chapter only discusses QoS.)

You can configure standard priority queueing and policing for different class maps within the same policy map. See the [“How QoS Features Interact” section on page 24-4](#) for information about valid QoS configurations.

To create a policy map, perform the following steps:

---

**Step 1** To add or edit a policy map, enter the following command:

```
hostname(config)# policy-map name
```

For example:

```
hostname(config)# policy-map QoS_policy
```

**Step 2** To configure priority queueing, enter the following commands:

```
hostname(config-pmap)# class priority_map_name
hostname(config-pmap-c)# priority
```

where the *priority\_map\_name* is the class map you created for prioritized traffic in “[Identifying Traffic for QoS Using Class Maps](#)” section on page 24-8.

For example:

```
hostname(config)# class-map priority-class
hostname(config-cmap)# match tunnel-group Tunnel-Group-1
hostname(config-cmap)# match dscp ef

hostname(config-cmap)# policy-map QoS_policy

hostname(config-pmap)# class priority_class
hostname(config-pmap-c)# priority
```

**Step 3** To configure policing, enter the following commands:

```
hostname(config-pmap)# class policing_map_name
hostname(config-pmap-c)# police {output | input} conform-rate [conform-burst]
[conform-action [drop | transmit]] [exceed-action [drop | transmit]]
```

where the *policing\_map\_name* is the class map you created for prioritized traffic in “[Identifying Traffic for QoS Using Class Maps](#)” section on page 24-8.

The *conform-burst* argument specifies the maximum number of instantaneous bytes allowed in a sustained burst before throttling to the conforming rate value, between 1000 and 512000000 bytes.

The **conform-action** keyword sets the action to take when the rate is less than the *conform-burst* value.

The *conform-rate* argument sets the rate limit for this traffic flow; between 8000 and 2000000000 bits per second.

The **drop** keyword drops the packet.

The **exceed-action** keyword sets the action to take when the rate is between the *conform-rate* value and the *conform-burst* value.

The **input** keyword enables policing of traffic flowing in the input direction.

The **output** keyword enables policing of traffic flowing in the output direction.

The **transmit** keyword transmits the packet.

For example:

```
hostname(config)# class-map policing-class
hostname(config-cmap)# match any

hostname(config-cmap)# policy-map QoS_policy

hostname(config-pmap)# class police_class
hostname(config-pmap-c)# police output 56000 10500
```

**Step 4** To activate the policy map on one or more interfaces, enter the following command:

```
hostname(config)# service-policy polycymap_name {global | interface interface_name}
```

Where **global** applies the policy map to all interfaces, and **interface** applies the policy to one interface. Only one global policy is allowed. Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with inspections, and an interface

policy with TCP normalization, then both inspections and TCP normalization are applied to the interface. However, if you have a global policy with inspections, and an interface policy with inspections, then only the interface policy inspections are applied to that interface.

In this example, the maximum rate for traffic of the `tcp_traffic` class is 56,000 bits/second and a maximum burst size of 10,500 bytes per second. For the `TC1-BestEffort` class, the maximum rate is 200,000 bits/second, with a maximum burst of 37,500 bytes/second. Traffic in the `TC1-voice` class has no policed maximum speed or burst rate because it belongs to a priority class.

```
hostname(config)# access-list tcp_traffic permit tcp any any
hostname(config)# class-map tcp_traffic
hostname(config-cmap)# match access-list tcp_traffic

hostname(config)# class-map TG1-voice
hostname(config-cmap)# match tunnel-group tunnel-grp1
hostname(config-cmap)# match dscp ef

hostname(config-cmap)# class-map TG1-BestEffort
hostname(config-cmap)# match tunnel-group tunnel-grp1
hostname(config-cmap)# match flow ip destination-address

hostname(config)# policy-map qos
hostname(config-pmap)# class tcp_traffic
hostname(config-pmap-c)# police output 56000 10500

hostname(config-pmap-c)# class TG1-voice
hostname(config-pmap-c)# priority

hostname(config-pmap-c)# class TG1-best-effort
hostname(config-pmap-c)# police output 200000 37500

hostname(config-pmap-c)# class class-default
hostname(config-pmap-c)# police output 1000000 37500

hostname(config-pmap-c)# service-policy qos global
```

## Creating a Policy for Traffic Shaping and Hierarchical Priority Queueing

You can create a policy map for an interface or globally for all interfaces that assigns QoS actions (and other feature actions) to the traffic in the class map. (See the [Chapter 21, “Using Modular Policy Framework,”](#) for information about other features. This chapter only discusses QoS.)

You can configure traffic shaping for all traffic on an interface, and optionally hierarchical priority queueing for a subset of latency-sensitive traffic. See the [“How QoS Features Interact”](#) section on [page 24-4](#) for information about valid QoS configurations.

If you want to configure hierarchical priority queueing, then first identify the traffic in [“Identifying Traffic for QoS Using Class Maps”](#) section on [page 24-8](#); traffic shaping always uses the `class-default` class map, which is automatically available.

**Note**

One side-effect of priority queueing is packet re-ordering. For IPSec packets, out-of-order packets that are not within the anti-replay window generate warning syslog messages. These warnings are false alarms in the case of priority queueing. You can configure the IPSec anti-replay window size to avoid possible false alarms. See the **crypto ipsec security-association replay** command in the *Cisco Security Appliance Command Reference*.

To create a policy map, perform the following steps:

- Step 1** (Optional) For hierarchical priority queueing, create a policy map that applies the priority queueing action to a class map by entering the following commands:

```
hostname(config)# policy-map name
hostname(config-pmap)# class priority_map_name
hostname(config-pmap-c)# priority
```

where the *priority\_map\_name* is the class map you created for prioritized traffic in [“Identifying Traffic for QoS Using Class Maps”](#) section on page 24-8.

For example:

```
hostname(config)# policy-map priority-sub-policy
hostname(config-pmap)# class priority-sub-map
hostname(config-pmap-c)# priority
```

- Step 2** To add or edit a policy map for traffic shaping, enter the following command:

```
hostname(config)# policy-map name
```

For example:

```
hostname(config)# policy-map shape_policy
```

- Step 3** To configure traffic shaping, enter the following commands:

```
hostname(config-pmap)# class class-default
hostname(config-pmap-c)# shape average rate [burst_size]
```

where the **average rate** argument sets the average rate of traffic in bits per second over a given fixed time period, between 64000 and 154400000. Specify a value that is a multiple of 8000. See the [“Traffic Shaping Overview”](#) section on page 24-4 for more information about how the time period is calculated.

The *burst\_size* argument sets the average burst size in bits that can be transmitted over a given fixed time period, between 2048 and 154400000. Specify a value that is a multiple of 128. If you do not specify the *burst\_size*, the default value is equivalent to 4-milliseconds of traffic at the specified average rate. For example, if the average rate is 1000000 bits per second, 4 ms worth =  $1000000 * 4/1000 = 4000$ .

You can only identify the **class-default** class map, which is defined as **match any**, because the security appliance requires all traffic to be matched for traffic shaping.

- Step 4** (Optional) To configure hierarchical priority queueing, enter the following command:

```
hostname(config-pmap-c)# service-policy priority_policy_map_name
```

where the *priority\_policy\_map\_name* is the policy map you created for prioritized traffic in [Step 1](#).

For example:

```
hostname(config)# policy-map priority-sub-policy
hostname(config-pmap)# class priority-sub-map
hostname(config-pmap-c)# priority
```

```
hostname(config-pmap-c)# policy-map shape_policy
hostname(config-pmap)# class class-default
hostname(config-pmap-c)# shape
hostname(config-pmap-c)# service-policy priority-sub-policy
```

**Step 5** To activate the policy map on an interface, enter the following command:

```
hostname(config)# service-policy polycymap_name interface interface_name
```



**Note**

You cannot configure traffic shaping in the global policy.

The following example enables traffic shaping on the outside interface, and limits traffic to 2 Mbps; priority queueing is enabled for VoIP traffic that is tagged with DSCP EF and AF13 and for IKE traffic:

```
hostname(config)# access-list ike permit udp any any eq 500
hostname(config)# class-map ike
hostname(config-cmap)# match access-list ike

hostname(config-cmap)# class-map voice_traffic
hostname(config-cmap)# match dscp EF AF13

hostname(config-cmap)# policy-map qos_class_policy
hostname(config-pmap)# class voice_traffic
hostname(config-pmap-c)# priority
hostname(config-pmap-c)# class ike
hostname(config-pmap-c)# priority

hostname(config-pmap-c)# policy-map qos_outside_policy
hostname(config-pmap)# class class-default
hostname(config-pmap-c)# shape average 2000000 16000
hostname(config-pmap-c)# service-policy qos_class_policy

hostname(config-pmap-c)# service-policy qos_outside_policy interface outside
```

## Viewing QoS Statistics

This section includes the following topics:

- [Viewing QoS Police Statistics, page 24-13](#)
- [Viewing QoS Standard Priority Statistics, page 24-14](#)
- [Viewing QoS Shaping Statistics, page 24-14](#)
- [Viewing QoS Standard Priority Queue Statistics, page 24-15](#)

## Viewing QoS Police Statistics

To view the QoS statistics for traffic policing, use the **show service-policy** command with the **police** keyword:

```
hostname# show service-policy police
```

The following is sample output for the **show service-policy police** command:

```

hostname# show service-policy police

Global policy:
 Service-policy: global_fw_policy

Interface outside:
 Service-policy: qos
 Class-map: browse
 police Interface outside:
 cir 56000 bps, bc 10500 bytes
 conformed 10065 packets, 12621510 bytes; actions: transmit
 exceeded 499 packets, 625146 bytes; actions: drop
 conformed 5600 bps, exceed 5016 bps
 Class-map: cmap2
 police Interface outside:
 cir 200000 bps, bc 37500 bytes
 conformed 17179 packets, 20614800 bytes; actions: transmit
 exceeded 617 packets, 770718 bytes; actions: drop
 conformed 198785 bps, exceed 2303 bps

```

## Viewing QoS Standard Priority Statistics

To view statistics for service policies implementing the **priority** command, use the **show service-policy** command with the **priority** keyword:

```
hostname# show service-policy priority
```

The following is sample output for the **show service-policy priority** command:

```

hostname# show service-policy priority
Global policy:
 Service-policy: global_fw_policy
Interface outside:
 Service-policy: qos
 Class-map: TG1-voice
 Priority:
 Interface outside: aggregate drop 0, aggregate transmit 9383

```



### Note

“Aggregate drop” denotes the aggregated drop in this interface; “aggregate transmit” denotes the aggregated number of transmitted packets in this interface.

## Viewing QoS Shaping Statistics

To view statistics for service policies implementing the **shape** command, use the **show service-policy** command with the **shape** keyword:

```
hostname# show service-policy shape
```

The following is sample output for the **show service-policy shape** command:

```

hostname# show service-policy shape
Interface outside
 Service-policy: shape
 Class-map: class-default

 Queueing
 queue limit 64 packets

```



```
(queue depth/total drops/no-buffer drops) 0/0/0
(pkts output/bytes output) 0/0

shape (average) cir 2000000, bc 8000, be 8000
```

The following is sample output of the **show service policy shape** command, which includes service policies that include the **shape** command and the **service-policy** command that calls the hierarchical priority policy and the related statistics:

```
hostname# show service-policy shape

Interface outside:
 Service-policy: shape
 Class-map: class-default

 Queueing
 queue limit 64 packets
 (queue depth/total drops/no-buffer drops) 0/0/0
 (pkts output/bytes output) 0/0

 shape (average) cir 2000000, bc 16000, be 16000

 Service-policy: voip
 Class-map: voip

 Queueing
 queue limit 64 packets
 (queue depth/total drops/no-buffer drops) 0/0/0
 (pkts output/bytes output) 0/0
 Class-map: class-default

 queue limit 64 packets
 (queue depth/total drops/no-buffer drops) 0/0/0
 (pkts output/bytes output) 0/0
```

## Viewing QoS Standard Priority Queue Statistics

To display the priority-queue statistics for an interface, use the **show priority-queue statistics** command in privileged EXEC mode. The results show the statistics for both the best-effort (BE) queue and the low-latency queue (LLQ). The following example shows the use of the **show priority-queue statistics** command for the interface named test, and the command output.

```
hostname# show priority-queue statistics test

Priority-Queue Statistics interface test

Queue Type = BE
Packets Dropped = 0
Packets Transmit = 0
Packets Enqueued = 0
Current Q Length = 0
Max Q Length = 0

Queue Type = LLQ
Packets Dropped = 0
Packets Transmit = 0
Packets Enqueued = 0
Current Q Length = 0
Max Q Length = 0
hostname#
```

In this statistical report, the meaning of the line items is as follows:

- “Packets Dropped” denotes the overall number of packets that have been dropped in this queue.
- “Packets Transmit” denotes the overall number of packets that have been transmitted in this queue.
- “Packets Enqueued” denotes the overall number of packets that have been queued in this queue.
- “Current Q Length” denotes the current depth of this queue.
- “Max Q Length” denotes the maximum depth that ever occurred in this queue.



## CHAPTER 25

# Configuring Application Layer Protocol Inspection

---

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the security appliance to do a deep packet inspection instead of passing the packet through the fast path (see the [“Stateful Inspection Overview” section on page 1-4](#) for more information about the fast path). As a result, inspection engines can affect overall throughput.

Several common inspection engines are enabled on the security appliance by default, but you might need to enable others depending on your network. This chapter includes the following sections:

- [Inspection Engine Overview, page 25-2](#)
  - [When to Use Application Protocol Inspection, page 25-2](#)
  - [Inspection Limitations, page 25-2](#)
  - [Default Inspection Policy, page 25-3](#)
- [Configuring Application Inspection, page 25-5](#)
- [CTIQBE Inspection, page 25-9](#)
- [DCERPC Inspection, page 25-11](#)
- [DNS Inspection, page 25-13](#)
- [ESMTP Inspection, page 25-23](#)
- [FTP Inspection, page 25-26](#)
- [GTP Inspection, page 25-32](#)
- [H.323 Inspection, page 25-38](#)
- [HTTP Inspection, page 25-44](#)
- [Instant Messaging Inspection, page 25-49](#)
- [ICMP Inspection, page 25-52](#)
- [ICMP Error Inspection, page 25-52](#)
- [ILS Inspection, page 25-53](#)
- [IPSec Pass Through Inspection, page 25-54](#)
- [MGCP Inspection, page 25-56](#)
- [NetBIOS Inspection, page 25-60](#)

- [PPTP Inspection, page 25-62](#)
- [RADIUS Accounting Inspection, page 25-62](#)
- [RSH Inspection, page 25-63](#)
- [RTSP Inspection, page 25-63](#)
- [SIP Inspection, page 25-65](#)
- [Skinny \(SCCP\) Inspection, page 25-71](#)
- [SMTP and Extended SMTP Inspection, page 25-74](#)
- [SNMP Inspection, page 25-76](#)
- [SQL\\*Net Inspection, page 25-76](#)
- [Sun RPC Inspection, page 25-77](#)
- [TFTP Inspection, page 25-79](#)
- [XDMCP Inspection, page 25-80](#)

## Inspection Engine Overview

This section includes the following topics:

- [When to Use Application Protocol Inspection, page 25-2](#)
- [Inspection Limitations, page 25-2](#)
- [Default Inspection Policy, page 25-3](#)

## When to Use Application Protocol Inspection

When a user establishes a connection, the security appliance checks the packet against access lists, creates an address translation, and creates an entry for the session in the fast path, so that further packets can bypass time-consuming checks. However, the fast path relies on predictable port numbers and does not perform address translations inside a packet.

Many protocols open secondary TCP or UDP ports. The initial session on a well-known port is used to negotiate dynamically assigned port numbers.

Other applications embed an IP address in the packet that needs to match the source address that is normally translated when it goes through the security appliance.

If you use applications like these, then you need to enable application inspection.

When you enable application inspection for a service that embeds IP addresses, the security appliance translates embedded addresses and updates any checksum or other fields that are affected by the translation.

When you enable application inspection for a service that uses dynamically assigned ports, the security appliance monitors sessions to identify the dynamic port assignments, and permits data exchange on these ports for the duration of the specific session.

## Inspection Limitations

See the following limitations for application protocol inspection:

- State information for multimedia sessions that require inspection are not passed over the state link for stateful failover. The exception is GTP, which is replicated over the state link.
- Some inspection engines do not support PAT, NAT, outside NAT, or NAT between same security interfaces. See [“Default Inspection Policy”](#) for more information about NAT support.

## Default Inspection Policy

By default, the configuration includes a policy that matches all default application inspection traffic and applies inspection to the traffic on all interfaces (a global policy). Default application inspection traffic includes traffic to the default ports for each protocol. You can only apply one global policy, so if you want to alter the global policy, for example, to apply inspection to non-standard ports, or to add inspections that are not enabled by default, you need to either edit the default policy or disable it and apply a new one.

[Table 25-1](#) lists all inspections supported, the default ports used in the default class map, and the inspection engines that are on by default, shown in bold. This table also notes any NAT limitations.

**Table 25-1 Supported Application Inspection Engines**

| Application <sup>1</sup>           | Default Port                                | NAT Limitations                                               | Standards <sup>2</sup>                   | Comments                                                                                                                                                              |
|------------------------------------|---------------------------------------------|---------------------------------------------------------------|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CTIQBE                             | TCP/2748                                    | —                                                             | —                                        | —                                                                                                                                                                     |
| DNS over UDP                       | UDP/53                                      | No NAT support is available for name resolution through WINS. | RFC 1123                                 | No PTR records are changed.                                                                                                                                           |
| <b>FTP</b>                         | TCP/21                                      | —                                                             | RFC 959                                  | —                                                                                                                                                                     |
| GTP                                | UDP/3386<br>UDP/2123                        | —                                                             | —                                        | Requires a special license.                                                                                                                                           |
| <b>H.323 H.225 and RAS</b>         | TCP/1720<br>UDP/1718<br>UDP (RAS) 1718-1719 | No NAT on same security interfaces.<br>No static PAT.         | ITU-T H.323, H.245, H225.0, Q.931, Q.932 | —                                                                                                                                                                     |
| HTTP                               | TCP/80                                      | —                                                             | RFC 2616                                 | Beware of MTU limitations stripping ActiveX and Java. If the MTU is too small to allow the Java or ActiveX tag to be included in one packet, stripping may not occur. |
| ICMP                               | —                                           | —                                                             | —                                        | All ICMP traffic is matched in the default class map.                                                                                                                 |
| ICMP ERROR                         | —                                           | —                                                             | —                                        | All ICMP traffic is matched in the default class map.                                                                                                                 |
| ILS (LDAP)                         | TCP/389                                     | No PAT.                                                       | —                                        | —                                                                                                                                                                     |
| MGCP                               | UDP/2427, 2727                              | —                                                             | RFC 2705bis-05                           | —                                                                                                                                                                     |
| <b>NetBIOS Name Server over IP</b> | UDP/137, 138 (Source ports)                 | —                                                             | —                                        | NetBIOS is supported by performing NAT of the packets for NBNS UDP port 137 and NBDS UDP port 138.                                                                    |
| PPTP                               | TCP/1723                                    | —                                                             | RFC 2637                                 | —                                                                                                                                                                     |

**Table 25-1 Supported Application Inspection Engines (continued)**

| Application <sup>1</sup>        | Default Port         | NAT Limitations                                        | Standards <sup>2</sup>           | Comments                                                                                                                                                                                                                                                           |
|---------------------------------|----------------------|--------------------------------------------------------|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RADIUS Accounting               | 1646                 | —                                                      | RFC 2865                         | —                                                                                                                                                                                                                                                                  |
| <b>RSH</b>                      | TCP/514              | No PAT                                                 | Berkeley UNIX                    | —                                                                                                                                                                                                                                                                  |
| RTSP                            | TCP/554              | No PAT.<br>No outside NAT.                             | RFC 2326, 2327, 1889             | No handling for HTTP cloaking.                                                                                                                                                                                                                                     |
| <b>SIP</b>                      | TCP/5060<br>UDP/5060 | No outside NAT.<br>No NAT on same security interfaces. | RFC 3261                         | —                                                                                                                                                                                                                                                                  |
| <b>SKINNY (SCCP)</b>            | TCP/2000             | No outside NAT.<br>No NAT on same security interfaces. | —                                | Does not handle TFTP uploaded Cisco IP Phone configurations under certain circumstances.                                                                                                                                                                           |
| <b>SMTP and ESMTP</b>           | TCP/25               | —                                                      | RFC 821, 1123, 1870              | —                                                                                                                                                                                                                                                                  |
| SNMP                            | UDP/161, 162         | No NAT or PAT.                                         | RFC 1155, 1157, 1212, 1213, 1215 | v.2 RFC 1902-1908; v.3 RFC 2570-2580.                                                                                                                                                                                                                              |
| <b>SQL*Net</b>                  | TCP/1521             | —                                                      | —                                | v.1 and v.2.                                                                                                                                                                                                                                                       |
| <b>Sun RPC over UDP and TCP</b> | UDP/111              | No NAT or PAT.                                         | —                                | The default class map includes UDP port 111; if you want to enable Sun RPC inspection for TCP port 111, you need to create a new class map that matches TCP port 111, add the class to the policy, and then apply the <b>inspect sunrpc</b> command to that class. |
| TFTP                            | UDP/69               | —                                                      | RFC 1350                         | Payload IP addresses are not translated.                                                                                                                                                                                                                           |
| <b>XDCMP</b>                    | UDP/177              | No NAT or PAT.                                         | —                                | —                                                                                                                                                                                                                                                                  |

1. Inspection engines that are enabled by default for the default port are in bold.
2. The security appliance is in compliance with these standards, but it does not enforce compliance on packets being inspected. For example, FTP commands are supposed to be in a particular order, but the security appliance does not enforce the order.

The default policy configuration includes the following commands:

```

class-map inspection_default
 match default-inspection-traffic
policy-map type inspect dns preset_dns_map
 parameters
 message-length maximum 512
policy-map global_policy
 class inspection_default
 inspect dns preset_dns_map
 inspect ftp
 inspect h323 h225
 inspect h323 ras
 inspect rsh
 inspect rtsp
 inspect esmtp
 inspect sqlnet
 inspect skinny

```

```
inspect sunrpc
inspect xdmcp
inspect sip
inspect netbios
inspect tftp
service-policy global_policy global
```

## Configuring Application Inspection

This feature uses Modular Policy Framework, so that implementing application inspection consists of identifying traffic, applying inspections to the traffic, and activating inspections on an interface. For some applications, you can perform special actions when you enable inspection. See [Chapter 21, “Using Modular Policy Framework,”](#) for more information.

Inspection is enabled by default for some applications. See the [“Default Inspection Policy”](#) section for more information. Use this section to modify your inspection policy.

To configure application inspection, perform the following steps:

- 
- Step 1** To identify the traffic to which you want to apply inspections, add either a Layer 3/4 class map for through traffic or a Layer 3/4 class map for management traffic. See the [“Creating a Layer 3/4 Class Map for Through Traffic”](#) section on page 21-5 and [“Creating a Layer 3/4 Class Map for Management Traffic”](#) section on page 21-7 for detailed information. The management Layer 3/4 class map can be used only with the RADIUS accounting inspection.

The default Layer 3/4 class map for through traffic is called “inspection\_default.” It matches traffic using a special **match** command, **match default-inspection-traffic**, to match the default ports for each application protocol.

You can specify a **match access-list** command along with the **match default-inspection-traffic** command to narrow the matched traffic to specific IP addresses. Because the **match default-inspection-traffic** command specifies the ports to match, any ports in the access list are ignored.

If you want to match non-standard ports, then create a new class map for the non-standard ports. See the [“Default Inspection Policy”](#) section on page 25-3 for the standard ports for each inspection engine. You can combine multiple class maps in the same policy if desired, so you can create one class map to match certain traffic, and another to match different traffic. However, if traffic matches a class map that contains an inspection command, and then matches another class map that also has an inspection command, only the first matching class is used. For example, SNMP matches the inspection\_default class. To enable SNMP inspection, enable SNMP inspection for the default class in [Step 5](#). Do not add another class that matches SNMP.

For example, to limit inspection to traffic from 10.1.1.0 to 192.168.1.0 using the default class map, enter the following commands:

```
hostname(config)# access-list inspect extended permit ip 10.1.1.0 255.255.255.0
192.168.1.0 255.255.255.0
hostname(config)# class-map inspection_default
hostname(config-cmap)# match access-list inspect
```

View the entire class map using the following command:

```
hostname(config-cmap)# show running-config class-map inspection_default
!
class-map inspection_default
 match default-inspection-traffic
 match access-list inspect
```

!

To inspect FTP traffic on port 21 as well as 1056 (a non-standard port), create an access list that specifies the ports, and assign it to a new class map:

```
hostname(config)# access-list ftp_inspect extended permit tcp any any eq 21
hostname(config)# access-list ftp_inspect extended permit tcp any any eq 1056
hostname(config)# class-map new_inspection
hostname(config-cmap)# match access-list ftp_inspect
```

**Step 2** (Optional) Some inspection engines let you control additional parameters when you apply the inspection to the traffic. See the following sections to configure an inspection policy map for your application:

- DCERPC—See the “Configuring a DCERPC Inspection Policy Map for Additional Inspection Control” section on page 25-12
- DNS—See the “Configuring a DNS Inspection Policy Map for Additional Inspection Control” section on page 25-20
- ESMTP—See the “Configuring an ESMTP Inspection Policy Map for Additional Inspection Control” section on page 25-24
- FTP—See the “Configuring an FTP Inspection Policy Map for Additional Inspection Control” section on page 25-28.
- GTP—See the “Configuring a GTP Inspection Policy Map for Additional Inspection Control” section on page 25-33.
- H323—See the “Configuring an H.323 Inspection Policy Map for Additional Inspection Control” section on page 25-40
- HTTP—See the “Configuring an HTTP Inspection Policy Map for Additional Inspection Control” section on page 25-45.
- Instant Messaging—See the “Configuring an Instant Messaging Inspection Policy Map for Additional Inspection Control” section on page 25-49
- MGCP—See the “Configuring an MGCP Inspection Policy Map for Additional Inspection Control” section on page 25-58.
- NetBIOS—See the “Configuring a NetBIOS Inspection Policy Map for Additional Inspection Control” section on page 25-60
- RADIUS Accounting—See the “Configuring a RADIUS Inspection Policy Map for Additional Inspection Control” section on page 25-63
- SIP—See the “Configuring a SIP Inspection Policy Map for Additional Inspection Control” section on page 25-66
- Skinny—See the “Configuring a Skinny (SCCP) Inspection Policy Map for Additional Inspection Control” section on page 25-73
- SNMP—See the “SNMP Inspection” section on page 25-76.

**Step 3** To add or edit a Layer 3/4 policy map that sets the actions to take with the class map traffic, enter the following command:

```
hostname(config)# policy-map name
hostname(config-pmap)#
```

The default policy map is called “global\_policy.” This policy map includes the default inspections listed in the “Default Inspection Policy” section on page 25-3. If you want to modify the default policy (for example, to add or delete an inspection, or to identify an additional class map for your actions), then enter **global\_policy** as the name.



- Step 4** To identify the class map from [Step 1](#) to which you want to assign an action, enter the following command:

```
hostname(config-pmap) # class class_map_name
hostname(config-pmap-c) #
```

If you are editing the default policy map, it includes the `inspection_default` class map. You can edit the actions for this class by entering **inspection\_default** as the name. To add an additional class map to this policy map, identify a different name. You can combine multiple class maps in the same policy if desired, so you can create one class map to match certain traffic, and another to match different traffic. However, if traffic matches a class map that contains an inspection command, and then matches another class map that also has an inspection command, only the first matching class is used. For example, SNMP matches the `inspection_default` class map. To enable SNMP inspection, enable SNMP inspection for the default class in [Step 5](#). Do not add another class that matches SNMP.

- Step 5** Enable application inspection by entering the following command:

```
hostname(config-pmap-c) # inspect protocol
```

The *protocol* is one of the following values:

**Table 25-2 Protocol Keywords**

| Keywords                                           | Notes                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ctiqbe</b>                                      | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>dcerpc</b> [ <i>map_name</i> ]                  | If you added a DCERPC inspection policy map according to “ <a href="#">Configuring a DCERPC Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-12, identify the map name in this command.                                                                                                                                                                                                                                                            |
| <b>dns</b> [ <i>map_name</i> ]                     | If you added a DNS inspection policy map according to “ <a href="#">Configuring a DNS Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-20, identify the map name in this command. The default DNS inspection policy map name is “ <code>preset_dns_map</code> .” The default inspection policy map sets the maximum DNS packet length to 512 bytes.                                                                                                |
| <b>esmtplib</b> [ <i>map_name</i> ]                | If you added an ESMTP inspection policy map according to “ <a href="#">Configuring an ESMTP Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-24, identify the map name in this command.                                                                                                                                                                                                                                                            |
| <b>ftplib</b> [ <b>strict</b> [ <i>map_name</i> ]] | Use the <b>strict</b> keyword to increase the security of protected networks by preventing web browsers from sending embedded commands in FTP requests. See the “ <a href="#">Using the strict Option</a> ” section on page 25-27 for more information.<br><br>If you added an FTP inspection policy map according to “ <a href="#">Configuring an FTP Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-28, identify the map name in this command. |
| <b>gtp</b> [ <i>map_name</i> ]                     | If you added a GTP inspection policy map according to the “ <a href="#">Configuring a GTP Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-33, identify the map name in this command.                                                                                                                                                                                                                                                              |

**Table 25-2 Protocol Keywords**

| Keywords                                     | Notes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>h323 h225</b> [ <i>map_name</i> ]         | If you added an H323 inspection policy map according to “ <a href="#">Configuring an H.323 Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-40, identify the map name in this command.                                                                                                                                                                                                                                                                          |
| <b>h323 ras</b> [ <i>map_name</i> ]          | If you added an H323 inspection policy map according to “ <a href="#">Configuring an H.323 Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-40, identify the map name in this command.                                                                                                                                                                                                                                                                          |
| <b>http</b> [ <i>map_name</i> ]              | If you added an HTTP inspection policy map according to the “ <a href="#">Configuring an HTTP Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-45, identify the map name in this command.                                                                                                                                                                                                                                                                       |
| <b>icmp</b>                                  | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>icmp error</b>                            | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>ils</b>                                   | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>im</b> [ <i>map_name</i> ]                | If you added an Instant Messaging inspection policy map according to “ <a href="#">Configuring an Instant Messaging Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-49, identify the map name in this command.                                                                                                                                                                                                                                                 |
| <b>mgcp</b> [ <i>map_name</i> ]              | If you added an MGCP inspection policy map according to “ <a href="#">Configuring an MGCP Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-58, identify the map name in this command.                                                                                                                                                                                                                                                                           |
| <b>netbios</b> [ <i>map_name</i> ]           | If you added a NetBIOS inspection policy map according to “ <a href="#">Configuring a NetBIOS Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-60, identify the map name in this command.                                                                                                                                                                                                                                                                       |
| <b>pptp</b>                                  | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>radius-accounting</b> [ <i>map_name</i> ] | The <b>radius-accounting</b> keyword is only available for a management class map. See the “ <a href="#">Creating a Layer 3/4 Class Map for Management Traffic</a> ” section on page 21-7 for more information about creating a management class map.<br><br>If you added a RADIUS accounting inspection policy map according to “ <a href="#">Configuring a RADIUS Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-63, identify the map name in this command. |
| <b>rsh</b>                                   | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>rtsp</b>                                  | —                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>sip</b> [ <i>map_name</i> ]               | If you added a SIP inspection policy map according to “ <a href="#">Configuring a SIP Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-66, identify the map name in this command.                                                                                                                                                                                                                                                                               |

**Table 25-2 Protocol Keywords**

| Keywords                          | Notes                                                                                                                                                                                                                                                              |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>skinny</b> [ <i>map_name</i> ] | If you added a Skinny inspection policy map according to “ <a href="#">Configuring a Skinny (SCCP) Inspection Policy Map for Additional Inspection Control</a> ” section on page 25-73, identify the map name in this command.                                     |
| <b>snmp</b> [ <i>map_name</i> ]   | If you added an SNMP inspection policy map according to “ <a href="#">SNMP Inspection</a> ” section on page 25-76, identify the map name in this command.                                                                                                          |
| <b>sqlnet</b>                     | —                                                                                                                                                                                                                                                                  |
| <b>sunrpc</b>                     | The default class map includes UDP port 111; if you want to enable Sun RPC inspection for TCP port 111, you need to create a new class map that matches TCP port 111, add the class to the policy, and then apply the <b>inspect sunrpc</b> command to that class. |
| <b>tftp</b>                       | —                                                                                                                                                                                                                                                                  |
| <b>xdmcp</b>                      | —                                                                                                                                                                                                                                                                  |

**Step 6** To activate the policy map on one or more interfaces, enter the following command:

```
hostname(config)# service-policy polycymap_name {global | interface interface_name}
```

Where **global** applies the policy map to all interfaces, and **interface** applies the policy to one interface. By default, the default policy map, “global\_policy,” is applied globally. Only one global policy is allowed. You can override the global policy on an interface by applying a service policy to that interface. You can only apply one policy map to each interface.

## CTIQBE Inspection

This section describes CTIQBE application inspection. This section includes the following topics:

- [CTIQBE Inspection Overview, page 25-9](#)
- [Limitations and Restrictions, page 25-10](#)
- [Verifying and Monitoring CTIQBE Inspection, page 25-10](#)

## CTIQBE Inspection Overview

CTIQBE protocol inspection supports NAT, PAT, and bidirectional NAT. This enables Cisco IP SoftPhone and other Cisco TAPI/JTAPI applications to work successfully with Cisco CallManager for call setup across the security appliance.

TAPI and JTAPI are used by many Cisco VoIP applications. CTIQBE is used by Cisco TSP to communicate with Cisco CallManager.

## Limitations and Restrictions

The following summarizes limitations that apply when using CTIQBE application inspection:

- CTIQBE application inspection does not support configurations with the **alias** command.
- Stateful failover of CTIQBE calls is not supported.
- Entering the **debug ctique** command may delay message transmission, which may have a performance impact in a real-time environment. When you enable this debugging or logging and Cisco IP SoftPhone seems unable to complete call setup through the security appliance, increase the timeout values in the Cisco TSP settings on the system running Cisco IP SoftPhone.

The following summarizes special considerations when using CTIQBE application inspection in specific scenarios:

- If two Cisco IP SoftPhones are registered with different Cisco CallManagers, which are connected to different interfaces of the security appliance, calls between these two phones fails.
- When Cisco CallManager is located on the higher security interface compared to Cisco IP SoftPhones, if NAT or outside NAT is required for the Cisco CallManager IP address, the mapping must be static as Cisco IP SoftPhone requires the Cisco CallManager IP address to be specified explicitly in its Cisco TSP configuration on the PC.
- When using PAT or Outside PAT, if the Cisco CallManager IP address is to be translated, its TCP port 2748 must be statically mapped to the same port of the PAT (interface) address for Cisco IP SoftPhone registrations to succeed. The CTIQBE listening port (TCP 2748) is fixed and is not user-configurable on Cisco CallManager, Cisco IP SoftPhone, or Cisco TSP.

## Verifying and Monitoring CTIQBE Inspection

The **show ctique** command displays information regarding the CTIQBE sessions established across the security appliance. It shows information about the media connections allocated by the CTIQBE inspection engine.

The following is sample output from the **show ctique** command under the following conditions. There is only one active CTIQBE session setup across the security appliance. It is established between an internal CTI device (for example, a Cisco IP SoftPhone) at local address 10.0.0.99 and an external Cisco CallManager at 172.29.1.77, where TCP port 2748 is the Cisco CallManager. The heartbeat interval for the session is 120 seconds.

```
hostname# # show ctique

Total: 1
 LOCAL FOREIGN STATE HEARTBEAT

1 10.0.0.99/1117 172.29.1.77/2748 1 120

RTP/RTCP: PAT xlates: mapped to 172.29.1.99(1028 - 1029)

MEDIA: Device ID 27 Call ID 0
 Foreign 172.29.1.99 (1028 - 1029)
 Local 172.29.1.88 (26822 - 26823)

```

The CTI device has already registered with the CallManager. The device internal address and RTP listening port is PATed to 172.29.1.99 UDP port 1028. Its RTCP listening port is PATed to UDP 1029.

The line beginning with `RTP/RTCP: PAT xlates:` appears only if an internal CTI device has registered with an external CallManager and the CTI device address and ports are PATED to that external interface. This line does not appear if the CallManager is located on an internal interface, or if the internal CTI device address and ports are translated to the same external interface that is used by the CallManager.

The output indicates a call has been established between this CTI device and another phone at 172.29.1.88. The RTP and RTCP listening ports of the other phone are UDP 26822 and 26823. The other phone locates on the same interface as the CallManager because the security appliance does not maintain a CTIQBE session record associated with the second phone and CallManager. The active call leg on the CTI device side can be identified with Device ID 27 and Call ID 0.

The following is sample output from the **show xlate debug** command for these CTIBQE connections:

```
hostname# show xlate debug
3 in use, 3 most used
Flags: D - DNS, d - dump, I - identity, i - inside, n - no random,
 r - portmap, s - static
TCP PAT from inside:10.0.0.99/1117 to outside:172.29.1.99/1025 flags ri idle 0:00:22
timeout 0:00:30
UDP PAT from inside:10.0.0.99/16908 to outside:172.29.1.99/1028 flags ri idle 0:00:00
timeout 0:04:10
UDP PAT from inside:10.0.0.99/16909 to outside:172.29.1.99/1029 flags ri idle 0:00:23
timeout 0:04:10
```

The **show conn state ctiqbe** command displays the status of CTIQBE connections. In the output, the media connections allocated by the CTIQBE inspection engine are denoted by a 'C' flag. The following is sample output from the **show conn state ctiqbe** command:

```
hostname# show conn state ctiqbe
1 in use, 10 most used
hostname# show conn state ctiqbe detail
1 in use, 10 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
 B - initial SYN from outside, C - CTIQBE media, D - DNS, d - dump,
 E - outside back connection, F - outside FIN, f - inside FIN,
 G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
 i - incomplete, J - GTP, j - GTP data, k - Skinny media,
 M - SMTP data, m - SIP media, O - outbound data, P - inside back connection,
 q - SQL*Net data, R - outside acknowledged FIN,
 R - UDP RPC, r - inside acknowledged FIN, S - awaiting inside SYN,
 s - awaiting outside SYN, T - SIP, t - SIP transient, U - up
```

## DCERPC Inspection

This section describes the DCERPC inspection engine. This section includes the following topics:

- [DCERPC Overview, page 25-11](#)
- [Configuring a DCERPC Inspection Policy Map for Additional Inspection Control, page 25-12](#)

## DCERPC Overview

DCERPC is a protocol widely used by Microsoft distributed client and server applications that allows software clients to execute programs on a server remotely.

This typically involves a client querying a server called the Endpoint Mapper listening on a well known port number for the dynamically allocated network information of a required service. The client then sets up a secondary connection to the server instance providing the service. The security appliance allows the appropriate port number and network address and also applies NAT, if needed, for the secondary connection.

DCERPC inspect maps inspect for native TCP communication between the EPM and client on well known TCP port 135. Map and lookup operations of the EPM are supported for clients. Client and server can be located in any security zone. The embedded server IP address and Port number are received from the applicable EPM response messages. Since a client may attempt multiple connections to the server port returned by EPM, multiple use of pinholes are allowed, which have user configurable timeouts.

## Configuring a DCERPC Inspection Policy Map for Additional Inspection Control

To specify additional DCERPC inspection parameters, create a DCERPC inspection policy map. You can then apply the inspection policy map when you enable DCERPC inspection according to the [“Configuring Application Inspection”](#) section on page 25-5.

To create a DCERPC inspection policy map, perform the following steps:

---

**Step 1** Create a DCERPC inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect dcerpc policy_map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

**Step 2** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

**Step 3** To configure parameters that affect the inspection engine, perform the following steps:

**a.** To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap)# parameters
hostname(config-pmap-p)#
```

**b.** To configure the timeout for DCERPC pinholes and override the global system pinhole timeout of two minutes, enter the following command:

```
hostname(config-pmap-p)# timeout pinhole hh:mm:ss
```

Where the *hh:mm:ss* argument is the timeout for pinhole connections. Value is between 0:0:1 and 1193:0:0.

**c.** To configure options for the endpoint mapper traffic, enter the following command:

```
hostname(config-pmap-p)# endpoint-mapper [service-only] [lookup-operation]
[timeout hh:mm:ss]
```

Where the *hh:mm:ss* argument is the timeout for pinholes generated from the lookup operation. If no timeout is configured for the lookup operation, the timeout pinhole command or the default is used. The **epm-service-only** keyword enforces endpoint mapper service during binding so that only its service traffic is processed. The **lookup-operation** keyword enables the lookup operation of the endpoint mapper service.

---

The following example shows how to define a DCERPC inspection policy map with the timeout configured for DCERPC pinholes.

```
hostname(config)# policy-map type inspect dcerpc dcerpc_map
hostname(config-pmap)# timeout pinhole 0:10:00

hostname(config)# class-map dcerpc
hostname(config-cmap)# match port tcp eq 135

hostname(config)# policy-map global-policy
hostname(config-pmap)# class dcerpc
hostname(config-pmap-c)# inspect msrpc dcerpc-map

hostname(config)# service-policy global-policy global
```

## DNS Inspection

This section describes DNS application inspection. This section includes the following topics:

- [How DNS Application Inspection Works, page 25-13](#)
- [How DNS Rewrite Works, page 25-14](#)
- [Configuring DNS Rewrite, page 25-15](#)
- [Verifying and Monitoring DNS Inspection, page 25-20](#)

## How DNS Application Inspection Works

The security appliance tears down the DNS session associated with a DNS query as soon as the DNS reply is forwarded by the security appliance. The security appliance also monitors the message exchange to ensure that the ID of the DNS reply matches the ID of the DNS query.

When DNS inspection is enabled, which is the default, the security appliance performs the following additional tasks:

- Translates the DNS record based on the configuration completed using the **alias**, **static** and **nat** commands (DNS Rewrite). Translation only applies to the A-record in the DNS reply; therefore, DNS Rewrite does not affect reverse lookups, which request the PTR record.




---

**Note** DNS Rewrite is not applicable for PAT because multiple PAT rules are applicable for each A-record and the PAT rule to use is ambiguous.

---

- Enforces the maximum DNS message length (the default is 512 bytes and the maximum length is 65535 bytes). The security appliance performs reassembly as needed to verify that the packet length is less than the maximum length configured. The security appliance drops the packet if it exceeds the maximum length.




---

**Note** If you enter the **inspect dns** command without the **maximum-length** option, DNS packet size is not checked

---

- Enforces a domain-name length of 255 bytes and a label length of 63 bytes.

- Verifies the integrity of the domain-name referred to by the pointer if compression pointers are encountered in the DNS message.
- Checks to see if a compression pointer loop exists.

A single connection is created for multiple DNS sessions, as long as they are between the same two hosts, and the sessions have the same 5-tuple (source/destination IP address, source/destination port, and protocol). DNS identification is tracked by *app\_id*, and the idle timer for each *app\_id* runs independently.

Because the *app\_id* expires independently, a legitimate DNS response can only pass through the security appliance within a limited period of time and there is no resource build-up. However, if you enter the **show conn** command, you will see the idle timer of a DNS connection being reset by a new DNS session. This is due to the nature of the shared DNS connection and is by design.

## How DNS Rewrite Works

When DNS inspection is enabled, DNS rewrite provides full support for NAT of DNS messages originating from any interface.

If a client on an inside network requests DNS resolution of an inside address from a DNS server on an outside interface, the DNS A-record is translated correctly. If the DNS inspection engine is disabled, the A-record is not translated.


As long as DNS inspection remains enabled, you can configure DNS rewrite using the **alias**, **static**, or **nat** commands. For details about the configuration required see the [“Configuring DNS Rewrite” section on page 25-15](#).

DNS Rewrite performs two functions:

- Translating a public address (the routable or “mapped” address) in a DNS reply to a private address (the “real” address) when the DNS client is on a private interface.
- Translating a private address to a public address when the DNS client is on the public interface.

In [Figure 25-1](#), the DNS server resides on the external (ISP) network. The real address of the server (192.168.100.1) has been mapped using the **static** command to the ISP-assigned address (209.165.200.5). When a web client on the inside interface attempts to access the web server with the URL <http://server.example.com>, the host running the web client sends a DNS request to the DNS server to resolve the IP address of the web server. The security appliance translates the non-routable source address in the IP header and forwards the request to the ISP network on its outside interface. When the DNS reply is returned, the security appliance applies address translation not only to the destination address, but also to the embedded IP address of the web server, which is contained in the A-record in the DNS reply. As a result, the web client on the inside network gets the correct address for connecting to the web server on the inside network. For configuration instructions for scenarios similar to this one, see the [“Configuring DNS Rewrite with Two NAT Zones” section on page 25-16](#).



**Figure 25-1** Translating the Address in a DNS Reply (DNS Rewrite)

DNS rewrite also works if the client making the DNS request is on a DMZ network and the DNS server is on an inside interface. For an illustration and configuration instructions for this scenario, see the [“DNS Rewrite with Three NAT Zones”](#) section on page 25-17.

## Configuring DNS Rewrite

You configure DNS rewrite using the **alias**, **static**, or **nat** commands. The **alias** and **static** command can be used interchangeably; however, we recommend using the **static** command for new deployments because it is more precise and unambiguous. Also, DNS rewrite is optional when using the **static** command.

This section describes how to use the **alias** and **static** commands to configure DNS rewrite. It provides configuration procedures for using the **static** command in a simple scenario and in a more complex scenario. Using the **nat** command is similar to using the **static** command except that DNS Rewrite is based on dynamic translation instead of a static mapping.

This section includes the following topics:

- [Using the Static Command for DNS Rewrite, page 25-15](#)
- [Using the Static Command for DNS Rewrite, page 25-15](#)
- [Configuring DNS Rewrite with Two NAT Zones, page 25-16](#)
- [DNS Rewrite with Three NAT Zones, page 25-17](#)
- [Configuring DNS Rewrite with Three NAT Zones, page 25-19](#)

For detailed syntax and additional functions for the **alias**, **nat**, and **static** command, see the appropriate command page in the *Cisco Security Appliance Command Reference*.

## Using the Static Command for DNS Rewrite

The **static** command causes addresses on an IP network residing on a specific interface to be translated into addresses on another IP network on a different interface. The syntax for this command is as follows:

```
hostname(config)# static (real_ifc,mapped_ifc) mapped-address real-address dns
```

The following example specifies that the address 192.168.100.10 on the inside interface is translated into 209.165.200.5 on the outside interface:

```
hostname(config)# static (inside,outside) 209.165.200.225 192.168.100.10 dns
```

**Note**

Using the **nat** command is similar to using the **static** command except that DNS Rewrite is based on dynamic translation instead of a static mapping.

## Using the Alias Command for DNS Rewrite

The **alias** command causes the security appliance to translate addresses on an IP network residing on any interface into addresses on another IP network connected through a different interface. The syntax for this command is as follows:

```
hostname(config)# alias (interface_name) mapped-address real-address
```

The following example specifies that the real address (192.168.100.10) on any interface except the inside interface will be translated to the mapped address (**209.165.200.225**) on the inside interface. Notice that the location of 192.168.100.10 is not precisely defined.

```
hostname(config)# alias (inside) 209.165.200.225 192.168.100.10
```

**Note**

If you use the **alias** command to configure DNS Rewrite, proxy ARP will be performed for the mapped address. To prevent this, disable Proxy ARP by entering the **sysopt noproxyarp** command after entering the **alias** command.

## Configuring DNS Rewrite with Two NAT Zones

To implement a DNS Rewrite scenario similar to the one shown in [Figure 25-1](#), perform the following steps:

**Step 1** Create a static translation for the web server, as follows:

```
hostname(config)# static (real_ifc,mapped_ifc) mapped-address real-address netmask
255.255.255.255 dns
```

where the arguments are as follows:

- *real\_ifc*—The name of the interface connected to the real addresses.
- *mapped\_ifc*—The name of the interface where you want the addresses to be mapped.
- *mapped-address*—The translated IP address of the web server.
- *real-address*—The real IP address of the web server.

**Step 2** Create an access list that permits traffic to the port that the web server listens to for HTTP requests.

```
hostname(config)# access-list acl-name extended permit tcp any host mapped-address eq port
```

where the arguments are as follows:

*acl-name*—The name you give the access list.  
*mapped-address*—The translated IP address of the web server.  
*port*—The TCP port that the web server listens to for HTTP requests.

**Step 3** Apply the access list created in [Step 2](#) to the mapped interface. To do so, use the **access-group** command, as follows:

```
hostname(config)# access-group acl-name in interface mapped_ifc
```

**Step 4** If DNS inspection is disabled or if you want to change the maximum DNS packet length, configure DNS inspection. DNS application inspection is enabled by default with a maximum DNS packet length of 512 bytes. For configuration instructions, see the [“Configuring Application Inspection” section on page 25-5](#).

**Step 5** On the public DNS server, add an A-record for the web server, such as:

```
domain-qualified-hostname. IN A mapped-address
```

where *domain-qualified-hostname* is the hostname with a domain suffix, as in `server.example.com`. The period after the hostname is important. *mapped-address* is the translated IP address of the web server.

The following example configures the security appliance for the scenario shown in [Figure 25-1](#). It assumes DNS inspection is already enabled.

```
hostname(config)# static (inside,outside) 209.165.200.225 192.168.100.1 netmask
255.255.255.255 dns
hostname(config)# access-list 101 permit tcp any host 209.165.200.225 eq www
hostname(config)# access-group 101 in interface outside
```


This configuration requires the following A-record on the DNS server:

```
server.example.com. IN A 209.165.200.225
```

## DNS Rewrite with Three NAT Zones

[Figure 25-2](#) provides a more complex scenario to illustrate how DNS inspection allows NAT to operate transparently with a DNS server with minimal configuration. For configuration instructions for scenarios like this one, see the [“Configuring DNS Rewrite with Three NAT Zones” section on page 25-19](#).

**Figure 25-2 DNS Rewrite with Three NAT Zones**



In [Figure 25-2](#), a web server, `server.example.com`, has the real address `192.168.100.10` on the DMZ interface of the security appliance. A web client with the IP address `10.10.10.25` is on the inside interface and a public DNS server is on the outside interface. The site NAT policies are as follows:

- The outside DNS server holds the authoritative address record for `server.example.com`.
- Hosts on the outside network can contact the web server with the domain name `server.example.com` through the outside DNS server or with the IP address `209.165.200.5`.
- Clients on the inside network can access the web server with the domain name `server.example.com` through the outside DNS server or with the IP address `192.168.100.10`.

When a host or client on any interface accesses the DMZ web server, it queries the public DNS server for the A-record of `server.example.com`. The DNS server returns the A-record showing that `server.example.com` binds to address `209.165.200.5`.

When a web client on the *outside* network attempts to access `http://server.example.com`, the sequence of events is as follows:

1. The host running the web client sends the DNS server a request for the IP address of `server.example.com`.
2. The DNS server responds with the IP address `209.165.200.225` in the reply.
3. The web client sends its HTTP request to `209.165.200.225`.
4. The packet from the outside host reaches the security appliance at the outside interface.
5. The static rule translates the address `209.165.200.225` to `192.168.100.10` and the security appliance directs the packet to the web server on the DMZ.

When a web client on the *inside* network attempts to access `http://server.example.com`, the sequence of events is as follows:

1. The host running the web client sends the DNS server a request for the IP address of `server.example.com`.
2. The DNS server responds with the IP address `209.165.200.225` in the reply.
3. The security appliance receives the DNS reply and submits it to the DNS application inspection engine.
4. The DNS application inspection engine does the following:

- a. Searches for any NAT rule to undo the translation of the embedded A-record address “[outside]:209.165.200.5”. In this example, it finds the following static configuration:

```
static (dmz,outside) 209.165.200.225 192.168.100.10 dns
```

- b. Uses the static rule to rewrite the A-record as follows because the **dns** option is included:

```
[outside]:209.165.200.225 --> [dmz]:192.168.100.10
```



**Note** If the **dns** option were not included with the **static** command, DNS Rewrite would not be performed and other processing for the packet continues.

- c. Searches for any NAT to translate the web server address, `[dmz]:192.168.100.10`, when communicating with the inside web client.

No NAT rule is applicable, so application inspection completes.

If a NAT rule (nat or static) were applicable, the **dns** option must also be specified. If the **dns** option were not specified, the A-record rewrite in step [b](#) would be reverted and other processing for the packet continues.

5. The security appliance sends the HTTP request to server.example.com on the DMZ interface.

## Configuring DNS Rewrite with Three NAT Zones

To enable the NAT policies for the scenario in [Figure 25-2](#), perform the following steps:

- 
- Step 1** Create a static translation for the web server on the DMZ network, as follows:
- ```
hostname(config)# static (dmz,outside) mapped-address real-address dns
```
- where the arguments are as follows:
- *dmz*—The name of the DMZ interface of the security appliance.
 - *outside*—The name of the outside interface of the security appliance.
 - *mapped-address*—The translated IP address of the web server.
 - *real-address*—The real IP address of the web server.
- Step 2** Create an access list that permits traffic to the port that the web server listens to for HTTP requests.
- ```
hostname(config)# access-list acl-name extended permit tcp any host mapped-address eq port
```
- where the arguments are as follows:
- acl-name*—The name you give the access list.
- mapped-address*—The translated IP address of the web server.
- port*—The TCP port that the web server listens to for HTTP requests.
- Step 3** Apply the access list created in [Step 2](#) to the outside interface. To do so, use the **access-group** command, as follows:
- ```
hostname(config)# access-group acl-name in interface outside
```
- Step 4** If DNS inspection is disabled or if you want to change the maximum DNS packet length, configure DNS inspection. DNS application inspection is enabled by default with a maximum DNS packet length of 512 bytes. For configuration instructions, see the “[Configuring Application Inspection](#)” section on [page 25-5](#).
- Step 5** On the public DNS server, add an A-record for the web server, such as:
- ```
domain-qualified-hostname. IN A mapped-address
```
- where *domain-qualified-hostname* is the hostname with a domain suffix, as in server.example.com. The period after the hostname is important. *mapped-address* is the translated IP address of the web server.
- 

The following example configures the security appliance for the scenario shown in [Figure 25-2](#). It assumes DNS inspection is already enabled.

```
hostname(config)# static (dmz,outside) 209.165.200.225 192.168.100.10 dns
hostname(config)# access-list 101 permit tcp any host 209.165.200.225 eq www
hostname(config)# access-group 101 in interface outside
```

This configuration requires the following A-record on the DNS server:

```
server.example.com. IN A 209.165.200.225
```

## Verifying and Monitoring DNS Inspection

To view information about the current DNS connections, enter the following command:

```
hostname# show conn
```

For connections using a DNS server, the source port of the connection may be replaced by the IP address of DNS server in the show conn command output.

A single connection is created for multiple DNS sessions, as long as they are between the same two hosts, and the sessions have the same 5-tuple (source/destination IP address, source/destination port, and protocol). DNS identification is tracked by app\_id, and the idle timer for each app\_id runs independently.

Because the app\_id expires independently, a legitimate DNS response can only pass through the security appliance within a limited period of time and there is no resource build-up. However, when you enter the **show conn** command, you see the idle timer of a DNS connection being reset by a new DNS session. This is due to the nature of the shared DNS connection and is by design.

To display the statistics for DNS application inspection, enter the **show service-policy** command. The following is sample output from the **show service-policy** command:

```
hostname# show service-policy
Interface outside:
 Service-policy: sample_policy
 Class-map: dns_port
 Inspect: dns maximum-length 1500, packet 0, drop 0, reset-drop 0
```

## Configuring a DNS Inspection Policy Map for Additional Inspection Control

DNS application inspection supports DNS message controls that provide protection against DNS spoofing and cache poisoning. User configurable rules allow filtering based on DNS header, domain name, resource record type and class. Zone transfer can be restricted between servers with this function, for example.

The Recursion Desired and Recursion Available flags in the DNS header can be masked to protect a public server from attack if that server only supports a particular internal zone. In addition, DNS randomization can be enabled to avoid spoofing and cache poisoning of servers that either do not support randomization, or utilize a weak pseudo random number generator. Limiting the domain names that can be queried also restricts the domain names which can be queried, which protects the public server further.

A configurable DNS mismatch alert can be used as notification if an excessive number of mismatching DNS responses are received, which could indicate a cache poisoning attack. In addition, a configurable check to enforce a Transaction Signature be attached to all DNS messages is also supported.

To specify actions when a message violates a parameter, create a DNS inspection policy map. You can then apply the inspection policy map when you enable DNS inspection according to the [“Configuring Application Inspection”](#) section on page 25-5.

To create a DNS inspection policy map, perform the following steps:

- 
- |               |                                                                                                                                                                                                                                                                                    |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | (Optional) Add one or more regular expressions for use in traffic matching commands according to the <a href="#">“Creating a Regular Expression”</a> section on page 21-12. See the types of text you can match in the <b>match</b> commands described in <a href="#">Step 3</a> . |
| <b>Step 2</b> | (Optional) Create one or more regular expression class maps to group regular expressions according to the <a href="#">“Creating a Regular Expression Class Map”</a> section on page 21-14.                                                                                         |

**Step 3** (Optional) Create a DNS inspection class map by performing the following steps.

A class map groups multiple traffic matches. Traffic must match *all* of the **match** commands to match the class map. You can alternatively identify **match** commands directly in the policy map. The difference between creating a class map and defining the traffic match directly in the inspection policy map is that the class map lets you create more complex match criteria, and you can reuse class maps.

To specify traffic that should not match the class map, use the **match not** command. For example, if the **match not** command specifies the string “example.com,” then any traffic that includes “example.com” does not match the class map.

For the traffic that you identify in this class map, you can specify actions such as drop, drop-connection, reset, mask, set the rate limit, and/or log the connection in the inspection policy map.

If you want to perform different actions for each **match** command, you should identify the traffic directly in the policy map.

- a. Create the class map by entering the following command:

```
hostname(config)# class-map type inspect dns [match-all] class_map_name
hostname(config-cmap)#
```

Where *class\_map\_name* is the name of the class map. The **match-all** keyword specifies that traffic must match all criteria to match the class map. **match-all** is the default and only option. The CLI enters class-map configuration mode, where you can enter one or more **match** commands.

- b. (Optional) To add a description to the class map, enter the following command:

```
hostname(config-cmap)# description string
```

- c. (Optional) To match a specific flag that is set in the DNS header, enter the following command:

```
hostname(config-cmap)# match [not] header-flag [eq] {f_well_known | f_value}
```

Where the *f\_well\_known* argument is the DNS flag bit. The *f\_value* argument is the 16-bit value in hex. The **eq** keyword specifies an exact match.

- d. (Optional) To match a DNS type, including Query type and RR type, enter the following command:

```
hostname(config-cmap)# match [not] dns-type {eq t_well_known | t_val} {range t_val1 t_val2}
```

Where the *t\_well\_known* argument is the DNS flag bit. The *t\_val* arguments are arbitrary values in the DNS type field (0-65535). The **range** keyword specifies a range and the **eq** keyword specifies an exact match.

- e. (Optional) To match a DNS class, enter the following command:

```
hostname(config-cmap)# match [not] dns-class {eq c_well_known | c_val} {range c_val1 c_val2}
```

Where the *c\_well\_known* argument is the DNS class. The *c\_val* arguments are arbitrary values in the DNS class field. The **range** keyword specifies a range and the **eq** keyword specifies an exact match.

- f. (Optional) To match a DNS question or resource record, enter the following command:

```
hostname(config-cmap)# match {question | {resource-record answer | authority | any}}
```

Where the **question** keyword specifies the question portion of a DNS message. The **resource-record** keyword specifies the resource record portion of a DNS message. The **answer** keyword specifies the Answer RR section. The **authority** keyword specifies the Authority RR section. The **additional** keyword specifies the Additional RR section.

- g. (Optional) To match a DNS message domain name list, enter the following command:

```
hostname(config-cmap)# match [not] domain-name {regex regex_id | regex class class_id}
```

The **regex** *regex\_name* argument is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#).

**Step 4** Create a DNS inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect dns policy_map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

**Step 5** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

**Step 6** To apply actions to matching traffic, perform the following steps.

a. Specify the traffic on which you want to perform actions using one of the following methods:

- Specify the DNS class map that you created in [Step 3](#) by entering the following command:

```
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

- Specify traffic directly in the policy map using one of the **match** commands described in [Step 3](#). If you use a **match not** command, then any traffic that does not match the criterion in the **match not** command has the action applied.

b. Specify the action you want to perform on the matching traffic by entering the following command:

```
hostname(config-pmap-c)# {[drop [send-protocol-error] |
drop-connection [send-protocol-error] | mask | reset] [log] | rate-limit message_rate}
```

Not all options are available for each **match** or **class** command. See the CLI help or the *Cisco Security Appliance Command Reference* for the exact options available.

The **drop** keyword drops all packets that match.

The **send-protocol-error** keyword sends a protocol error message.

The **drop-connection** keyword drops the packet and closes the connection.

The **mask** keyword masks out the matching portion of the packet.

The **reset** keyword drops the packet, closes the connection, and sends a TCP reset to the server and/or client.

The **log** keyword, which you can use alone or with one of the other keywords, sends a system log message.

The **rate-limit** *message\_rate* argument limits the rate of messages.

You can specify multiple **class** or **match** commands in the policy map. For information about the order of **class** and **match** commands, see the “[Defining Actions in an Inspection Policy Map](#)” section on [page 21-8](#).

**Step 7** To configure parameters that affect the inspection engine, perform the following steps:

a. To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap)# parameters
hostname(config-pmap-p)#
```

b. To randomize the DNS identifier for a DNS query, enter the following command:

```
hostname(config-pmap-p)# id-randomization
```



- c. To enable logging for excessive DNS ID mismatches, enter the following command:

```
hostname(config-pmap-p)# id-mismatch [count number duration seconds] action log
```

Where the **count** *string* argument specifies the maximum number of mismatch instances before a system message log is sent. The **duration seconds** specifies the period, in seconds, to monitor.

- d. To require a TSIG resource record to be present, enter the following command:

```
hostname(config-pmap-p)# tsig enforced action {drop [log] | [log]}
```

Where the **count** *string* argument specifies the maximum number of mismatch instances before a system message log is sent. The **duration seconds** specifies the period, in seconds, to monitor.

---

The following example shows a how to define a DNS inspection policy map.

```
hostname(config)# regex domain_example "example\.com"
hostname(config)# regex domain_foo "foo\.com"

hostname(config)# ! define the domain names that the server serves
hostname(config)# class-map type inspect regex match-any my_domains
hostname(config-cmap)# match regex domain_example
hostname(config-cmap)# match regex domain_foo

hostname(config)# ! Define a DNS map for query only
hostname(config)# class-map type inspect dns match-all pub_server_map
hostname(config-cmap)# match not header-flag QR
hostname(config-cmap)# match question
hostname(config-cmap)# match not domain-name regex class my_domains

hostname(config)# policy-map type inspect dns serv_prot
hostname(config-pmap)# class pub_server_map
hostname(config-pmap-c)# drop log
hostname(config-pmap-c)# match header-flag RD
hostname(config-pmap-c)# mask log

hostname(config)# class-map dns_serv_map
hostname(config-cmap)# match default-inspection-traffic

hostname(config)# policy-map pub_policy
hostname(config-pmap)# class dns_serv_map
hostname(config-pmap-c)# inspect dns serv_prot

hostname(config)# service-policy pub_policy interface dmz
```

## ESMTP Inspection

ESMTP inspection detects attacks, including spam, phishing, malformed message attacks, buffer overflow/underflow attacks. It also provides support for application security and protocol conformance, which enforce the sanity of the ESMTP messages as well as detect several attacks, block senders/receivers, and block mail relay.

## Configuring an ESMTP Inspection Policy Map for Additional Inspection Control

To specify actions when a message violates a parameter, create an ESMTP inspection policy map. You can then apply the inspection policy map when you enable ESMTP inspection according to the [“Configuring Application Inspection” section on page 25-5](#).

To create an ESMTP inspection policy map, perform the following steps:

**Step 1** (Optional) Add one or more regular expressions for use in traffic matching commands according to the [“Creating a Regular Expression” section on page 21-12](#). See the types of text you can match in the **match** commands described in [Step 3](#).

**Step 2** (Optional) Create one or more regular expression class maps to group regular expressions according to the [“Creating a Regular Expression Class Map” section on page 21-14](#).

**Step 3** Create an ESMTP inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect esmtp policy_map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

**Step 4** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

**Step 5** To apply actions to matching traffic, perform the following steps.

- a. Specify traffic directly in the policy map using one of the **match** commands described in [Step 3](#). If you use a **match not** command, then any traffic that does not match the criterion in the **match not** command has the action applied.
- b. Specify the action you want to perform on the matching traffic by entering the following command:

```
hostname(config-pmap-c)# {[drop [send-protocol-error] |
drop-connection [send-protocol-error] | mask | reset] [log] | rate-limit message_rate}
```

Not all options are available for each **match** or **class** command. See the CLI help or the *Cisco Security Appliance Command Reference* for the exact options available.

The **drop** keyword drops all packets that match.

The **send-protocol-error** keyword sends a protocol error message.

The **drop-connection** keyword drops the packet and closes the connection.

The **mask** keyword masks out the matching portion of the packet.

The **reset** keyword drops the packet, closes the connection, and sends a TCP reset to the server and/or client.

The **log** keyword, which you can use alone or with one of the other keywords, sends a system log message.

The **rate-limit** *message\_rate* argument limits the rate of messages.

You can specify multiple **class** or **match** commands in the policy map. For information about the order of **class** and **match** commands, see the [“Defining Actions in an Inspection Policy Map” section on page 21-8](#).

**Step 6** To configure parameters that affect the inspection engine, perform the following steps:

- a. To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap)# parameters
hostname(config-pmap-p)#
```

- b. To configure a local domain name, enter the following command:

```
hostname(config-pmap-p)# mail-relay domain-name action [drop-connection / log]
```

Where the **drop-connection** action closes the connection. The **log** action sends a system log message when this policy map matches traffic.

- c. To enforce banner obfuscation, enter the following command:

```
hostname(config-pmap-p)# mask-banner
```

- d. (Optional) To detect special characters in sender or receiver email addresses, enter the following command:

```
hostname(config-pmap-p)# special-character action [drop-connection | log]
```

Using this command detects pipe (|), backquote (`) and null characters.

- e. (Optional) To match the body length or body line length, enter the following command:

```
hostname(config-pmap-p)# match body [line] length gt length
```

Where *length* is the length of the message body or the length of a line in the message body.

- f. (Optional) To match an ESMTP command verb, enter the following command:

```
hostname(config-pmap-p)# match cmd verb verb
```

Where *verb* is any of the following ESMTP commands:

```
AUTH|DATA|EHLO|ETRN|HELO|HELP|MAIL|NOOP|QUIT|RCPT|RSET|SAML|SOML|VRFY
```

- g. (Optional) To match the number of recipient addresses, enter the following command:

```
hostname(config-pmap-p)# match cmd RCPT count gt count
```

Where *count* is the number of recipient addresses.

- h. (Optional) To match the command line length, enter the following command:

```
hostname(config-pmap-p)# match cmd line length gt length
```

Where *length* is the command line length.

- i. (Optional) To match the ehlo-reply-parameters, enter the following command:

```
hostname(config-pmap-p)# match ehlo-reply-parameter extensions
```

Where *extensions* are the ESMTP service extensions sent by the server in response to the EHLO message from the client. These extensions are implemented as a new command or as parameters to an existing command. *extensions* can be any of the following:

```
8bitmime|binarymime|checkpoint|dsn|ecode|etrn|others|pipelining|size|vrfy
```

- j. (Optional) To match the header length or header line length, enter the following command:

```
hostname(config-pmap-p)# match header [line] length gt length
```

Where *length* is the number of characters in the header or line.

- k. (Optional) To match the header to-fields count, enter the following command:

```
hostname(config-pmap-p)# match header to-fields count gt count
```

Where *count* is the number of recipients in the to-field of the header.

- l. (Optional) To match the number of invalid recipients, enter the following command:

```
hostname(config-pmap-p)# match invalid-recipients count gt count
```

Where *count* is the number of invalid recipients.

- m. (Optional) To match the type of MIME encoding scheme used, enter the following command:

```
hostname(config-pmap-p)# match mime encoding [7bit|8bit|base64|binary|others|
quoted-printable]
```

- n. (Optional) To match the MIME filename length, enter the following command:

```
hostname(config-pmap-p)# match mime filename length gt length
```

Where *length* is the length of the *filename* in the range 1 to 1000.

- o. (Optional) To match the MIME file type, enter the following command:

```
hostname(config-pmap-p)# match mime filetype regex [name | class name]
```

Where *name* or *class name* is the regular expression that matches a file type or a class map. The regular expression used to match a class map can select multiple file types.

- p. (Optional) To match a sender address, enter the following command:

```
hostname(config-pmap-p)# match sender-address regex [name | class name]
```

Where *name* or *class name* is the regular expression that matches a sender address or a class map. The regular expression used to match a class map can select multiple sender addresses.

- q. (Optional) To match the length of a sender's address, enter the following command:

```
hostname(config-pmap-p)# match sender-address length gt length
```

Where *length* is the number of characters in the sender's address.

The following example shows how to define an ESMTP inspection policy map.

```
hostname(config)# regex user1 "user1@cisco.com"
hostname(config)# regex user2 "user2@cisco.com"
hostname(config)# regex user3 "user3@cisco.com"
hostname(config)# class-map type regex senders_black_list
hostname(config-cmap)# description "Regular expressions to filter out undesired senders"
hostname(config-cmap)# match regex user1
hostname(config-cmap)# match regex user2
hostname(config-cmap)# match regex user3

hostname(config)# policy-map type inspect esmtp advanced_esmtp_map
hostname(config-pmap)# match sender-address regex class senders_black_list
hostname(config-pmap-c)# drop-connection log

hostname(config)# policy-map outside_policy
hostname(config-pmap)# class inspection_default
hostname(config-pmap-c)# inspect esmtp advanced_esmtp_map

hostname(config)# service-policy outside_policy interface outside
```

## FTP Inspection

This section describes the FTP inspection engine. This section includes the following topics:

- [FTP Inspection Overview, page 25-27](#)
- [Using the strict Option, page 25-27](#)
- [Configuring an FTP Inspection Policy Map for Additional Inspection Control, page 25-28](#)
- [Verifying and Monitoring FTP Inspection, page 25-31](#)

## FTP Inspection Overview

The FTP application inspection inspects the FTP sessions and performs four tasks:

- Prepares dynamic secondary data connection
- Tracks the FTP command-response sequence
- Generates an audit trail
- Translates the embedded IP address

FTP application inspection prepares secondary channels for FTP data transfer. Ports for these channels are negotiated through PORT or PASV commands. The channels are allocated in response to a file upload, a file download, or a directory listing event.



### Note

If you disable FTP inspection engines with the **no inspect ftp** command, outbound users can start connections only in passive mode, and all inbound FTP is disabled.

## Using the strict Option

Using the **strict** option with the **inspect ftp** command increases the security of protected networks by preventing web browsers from sending embedded commands in FTP requests.



### Note

To specify FTP commands that are not permitted to pass through the security appliance, create an FTP map according to the [“Configuring an FTP Inspection Policy Map for Additional Inspection Control” section on page 25-28](#).

After you enable the **strict** option on an interface, FTP inspection enforces the following behavior:

- An FTP command must be acknowledged before the security appliance allows a new command.
- The security appliance drops connections that send embedded commands.
- The 227 and PORT commands are checked to ensure they do not appear in an error string.



### Caution

Using the **strict** option may cause the failure of FTP clients that are not strictly compliant with FTP RFCs.

If the **strict** option is enabled, each FTP command and response sequence is tracked for the following anomalous activity:

- Truncated command—Number of commas in the PORT and PASV reply command is checked to see if it is five. If it is not five, then the PORT command is assumed to be truncated and the TCP connection is closed.

- **Incorrect command**—Checks the FTP command to see if it ends with <CR><LF> characters, as required by the RFC. If it does not, the connection is closed.
- **Size of RETR and STOR commands**—These are checked against a fixed constant. If the size is greater, then an error message is logged and the connection is closed.
- **Command spoofing**—The PORT command should always be sent from the client. The TCP connection is denied if a PORT command is sent from the server.
- **Reply spoofing**—PASV reply command (227) should always be sent from the server. The TCP connection is denied if a PASV reply command is sent from the client. This prevents the security hole when the user executes “227 xxxxx a1, a2, a3, a4, p1, p2.”
- **TCP stream editing**—The security appliance closes the connection if it detects TCP stream editing.
- **Invalid port negotiation**—The negotiated dynamic port value is checked to see if it is less than 1024. As port numbers in the range from 1 to 1024 are reserved for well-known connections, if the negotiated port falls in this range, then the TCP connection is freed.
- **Command pipelining**—The number of characters present after the port numbers in the PORT and PASV reply command is cross checked with a constant value of 8. If it is more than 8, then the TCP connection is closed.
- The security appliance replaces the FTP server response to the SYST command with a series of Xs. to prevent the server from revealing its system type to FTP clients. To override this default behavior, use the **no mask-syst-reply** command in the FTP map.

## Configuring an FTP Inspection Policy Map for Additional Inspection Control

FTP command filtering and security checks are provided using strict FTP inspection for improved security and control. Protocol conformance includes packet length checks, delimiters and packet format checks, command terminator checks, and command validation.

Blocking FTP based on user values is also supported so that it is possible for FTP sites to post files for download, but restrict access to certain users. You can block FTP connections based on file type, server name, and other attributes. System message logs are generated if an FTP connection is denied after inspection.

If you want FTP inspection to allow FTP servers to reveal their system type to FTP clients, and limit the allowed FTP commands, then create and configure an FTP map. You can then apply the FTP map when you enable FTP inspection according to the [“Configuring Application Inspection” section on page 25-5](#).

To create an FTP map, perform the following steps:

- 
- Step 1** (Optional) Add one or more regular expressions for use in traffic matching commands according to the [“Creating a Regular Expression” section on page 21-12](#). See the types of text you can match in the **match** commands described in [Step 3](#).
- Step 2** (Optional) Create one or more regular expression class maps to group regular expressions according to the [“Creating a Regular Expression Class Map” section on page 21-14](#).
- Step 3** (Optional) Create an FTP inspection class map by performing the following steps.

A class map groups multiple traffic matches. Traffic must match *all* of the **match** commands to match the class map. You can alternatively identify **match** commands directly in the policy map. The difference between creating a class map and defining the traffic match directly in the inspection policy map is that the class map lets you create more complex match criteria, and you can reuse class maps.

To specify traffic that should not match the class map, use the **match not** command. For example, if the **match not** command specifies the string “example.com,” then any traffic that includes “example.com” does not match the class map.

For the traffic that you identify in this class map, you can specify actions such as drop, drop-connection, reset, mask, set the rate limit, and/or log the connection in the inspection policy map.

If you want to perform different actions for each **match** command, you should identify the traffic directly in the policy map.

- a. Create the class map by entering the following command:

```
hostname(config)# class-map type inspect ftp [match-all] class_map_name
hostname(config-cmap)#
```

Where *class\_map\_name* is the name of the class map. The **match-all** keyword specifies that traffic must match all criteria to match the class map. **match-all** is the default and only option. The CLI enters class-map configuration mode, where you can enter one or more **match** commands.

- b. (Optional) To add a description to the class map, enter the following command:

```
hostname(config-cmap)# description string
```

- c. (Optional) To match a filename for FTP transfer, enter the following command:

```
hostname(config-cmap)# match [not] filename regex [regex_name |
class regex_class_name]
```

Where the *regex\_name* is the regular expression you created in [Step 1](#). The **class regex\_class\_name** is the regular expression class map you created in [Step 2](#).

- d. (Optional) To match a file type for FTP transfer,, enter the following command:

```
hostname(config-cmap)# match [not] filetype regex [regex_name |
class regex_class_name]
```

Where the *regex\_name* is the regular expression you created in [Step 1](#). The **class regex\_class\_name** is the regular expression class map you created in [Step 2](#).

- e. (Optional) To disallow specific FTP commands, use the following command:

```
hostname(config-cmap)# match [not] request-command ftp_command [ftp_command...]
```

Where *ftp\_command* with one or more FTP commands that you want to restrict. See [Table 25-3](#) for a list of the FTP commands that you can restrict.

**Table 25-3 FTP Map request-command deny Options**

| request-command deny Option | Purpose                                                                                      |
|-----------------------------|----------------------------------------------------------------------------------------------|
| <b>appe</b>                 | Disallows the command that appends to a file.                                                |
| <b>cdup</b>                 | Disallows the command that changes to the parent directory of the current working directory. |
| <b>dele</b>                 | Disallows the command that deletes a file on the server.                                     |
| <b>get</b>                  | Disallows the client command for retrieving a file from the server.                          |
| <b>help</b>                 | Disallows the command that provides help information.                                        |
| <b>mkd</b>                  | Disallows the command that makes a directory on the server.                                  |
| <b>put</b>                  | Disallows the client command for sending a file to the server.                               |
| <b>rmd</b>                  | Disallows the command that deletes a directory on the server.                                |

**Table 25-3** FTP Map request-command deny Options (continued)

| request-command deny Option | Purpose                                                                                               |
|-----------------------------|-------------------------------------------------------------------------------------------------------|
| <b>rnfr</b>                 | Disallows the command that specifies rename-from filename.                                            |
| <b>rnto</b>                 | Disallows the command that specifies rename-to filename.                                              |
| <b>site</b>                 | Disallows the command that are specific to the server system. Usually used for remote administration. |
| <b>stou</b>                 | Disallows the command that stores a file using a unique file name.                                    |

- f. (Optional) To match an FTP server, enter the following command:

```
hostname(config-cmap)# match [not] server regex [regex_name | class regex_class_name]
```

Where the *regex\_name* is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#).

- g. (Optional) To match an FTP username, enter the following command:

```
hostname(config-cmap)# match [not] username regex [regex_name |
class regex_class_name]
```

Where the *regex\_name* is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#).

- Step 4** Create an FTP inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect ftp policy_map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

- Step 5** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

- Step 6** To apply actions to matching traffic, perform the following steps.

- a. Specify the traffic on which you want to perform actions using one of the following methods:

- Specify the FTP class map that you created in [Step 3](#) by entering the following command:

```
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

- Specify traffic directly in the policy map using one of the **match** commands described in [Step 3](#). If you use a **match not** command, then any traffic that does not match the criterion in the **match not** command has the action applied.

- b. Specify the action you want to perform on the matching traffic by entering the following command:

```
hostname(config-pmap-c)# {[drop [send-protocol-error] |
drop-connection [send-protocol-error] | mask | reset] [log] | rate-limit message_rate}
```

Not all options are available for each **match** or **class** command. See the CLI help or the *Cisco Security Appliance Command Reference* for the exact options available.

The **drop** keyword drops all packets that match.

The **send-protocol-error** keyword sends a protocol error message.

The **drop-connection** keyword drops the packet and closes the connection.



The **mask** keyword masks out the matching portion of the packet.

The **reset** keyword drops the packet, closes the connection, and sends a TCP reset to the server and/or client.

The **log** keyword, which you can use alone or with one of the other keywords, sends a system log message.

The **rate-limit** *message\_rate* argument limits the rate of messages.

You can specify multiple **class** or **match** commands in the policy map. For information about the order of **class** and **match** commands, see the “[Defining Actions in an Inspection Policy Map](#)” section on page 21-8.

**Step 7** To configure parameters that affect the inspection engine, perform the following steps:

- a. To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap) # parameters
hostname(config-pmap-p) #
```

- b. To mask the greeting banner from the FTP server, enter the following command:

```
hostname(config-pmap-p) # mask-banner
```

- c. To mask the reply to **syst** command, enter the following command:

```
hostname(config-pmap-p) # mask-syst-reply
```

Before submitting a username and password, all FTP users are presented with a greeting banner. By default, this banner includes version information useful to hackers trying to identify weaknesses in a system. The following example shows how to mask this banner:

```
hostname(config) # policy-map type inspect ftp mymap
hostname(config-pmap) # parameters
hostname(config-pmap-p) # mask-banner

hostname(config) # class-map match-all ftp-traffic
hostname(config-cmap) # match port tcp eq ftp

hostname(config) # policy-map ftp-policy
hostname(config-pmap) # class ftp-traffic
hostname(config-pmap-c) # inspect ftp strict mymap

hostname(config) # service-policy ftp-policy interface inside
```

## Verifying and Monitoring FTP Inspection

FTP application inspection generates the following log messages:

- An Audit record 302002 is generated for each file that is retrieved or uploaded.
- The FTP command is checked to see if it is RETR or STOR and the retrieve and store commands are logged.
- The username is obtained by looking up a table providing the IP address.
- The username, source IP address, destination IP address, NAT address, and the file operation are logged.

- Audit record 201005 is generated if the secondary dynamic channel preparation failed due to memory shortage.

In conjunction with NAT, the FTP application inspection translates the IP address within the application payload. This is described in detail in RFC 959.

## GTP Inspection

This section describes the GTP inspection engine. This section includes the following topics:

- [GTP Inspection Overview, page 25-32](#)
- [Configuring a GTP Inspection Policy Map for Additional Inspection Control, page 25-33](#)
- [Verifying and Monitoring GTP Inspection, page 25-37](#)


**Note**

GTP inspection requires a special license. If you enter GTP-related commands on a security appliance without the required license, the security appliance displays an error message.

## GTP Inspection Overview

GPRS provides uninterrupted connectivity for mobile subscribers between GSM networks and corporate networks or the Internet. The GGSN is the interface between the GPRS wireless data network and other networks. The SGSN performs mobility, data session management, and data compression (See [Figure 25-3](#)).

**Figure 25-3** GPRS Tunneling Protocol



The UMTS is the commercial convergence of fixed-line telephony, mobile, Internet and computer technology. UTRAN is the networking protocol used for implementing wireless networks in this system. GTP allows multi-protocol packets to be tunneled through a UMTS/GPRS backbone between a GGSN, an SGSN and the UTRAN.

GTP does not include any inherent security or encryption of user data, but using GTP with the security appliance helps protect your network against these risks.

The SGSN is logically connected to a GGSN using GTP. GTP allows multiprotocol packets to be tunneled through the GPRS backbone between GSNs. GTP provides a tunnel control and management protocol that allows the SGSN to provide GPRS network access for a mobile station by creating, modifying, and deleting tunnels. GTP uses a tunneling mechanism to provide a service for carrying user data packets.

**Note**

When using GTP with failover, if a GTP connection is established and the active unit fails before data is transmitted over the tunnel, the GTP data connection (with a “j” flag set) is not replicated to the standby unit. This occurs because the active unit does not replicate embryonic connections to the standby unit.

## Configuring a GTP Inspection Policy Map for Additional Inspection Control

If you want to enforce additional parameters on GTP traffic, create and configure a GTP map. If you do not specify a map with the **inspect gtp** command, the security appliance uses the default GTP map, which is preconfigured with the following default values:

- **request-queue 200**
- **timeout gsn 0:30:00**
- **timeout pdp-context 0:30:00**
- **timeout request 0:01:00**
- **timeout signaling 0:30:00**
- **timeout tunnel 0:01:00**
- **tunnel-limit 500**

To create and configure a GTP map, perform the following steps. You can then apply the GTP map when you enable GTP inspection according to the [“Configuring Application Inspection” section on page 25-5](#).

**Step 1** Create a GTP inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect gtp policy_map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

**Step 2** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

**Step 3** To match an Access Point name, enter the following command:

```
hostname(config-pmap)# match [not] apn regex [regex_name | class regex_class_name]
```

Where the *regex\_name* is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#).

**Step 4** To match a message ID, enter the following command:

```
hostname(config-pmap)# match [not] message id [message_id | range lower_range upper_range]
```

Where the *message\_id* is an alphanumeric identifier between 1 and 255. The *lower\_range* is lower range of message IDs. The *upper\_range* is the upper range of message IDs.

- Step 5** To match a message length, enter the following command:

```
hostname(config-pmap)# match [not] message length min min_length max max_length
```

Where the *min\_length* and *max\_length* are both between 1 and 65536. The length specified by this command is the sum of the GTP header and the rest of the message, which is the payload of the UDP packet.

- Step 6** To match the version, enter the following command:

```
hostname(config-pmap)# match [not] version [version_id | range lower_range upper_range]
```

Where the *version\_id* is between 0 and 255. The *lower\_range* is lower range of versions. The *upper\_range* is the upper range of versions.

- Step 7** To configure parameters that affect the inspection engine, perform the following steps:

- a. To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap)# parameters
hostname(config-pmap-p)#
```

The **mnc** *network\_code* argument is a two or three-digit value identifying the network code.

By default, the security appliance does not check for valid MCC/MNC combinations. This command is used for IMSI Prefix filtering. The MCC and MNC in the IMSI of the received packet is compared with the MCC/MNC configured with this command and is dropped if it does not match.

This command must be used to enable IMSI Prefix filtering. You can configure multiple instances to specify permitted MCC and MNC combinations. By default, the security appliance does not check the validity of MNC and MCC combinations, so you must verify the validity of the combinations configured. To find more information about MCC and MNC codes, see the ITU E.212 recommendation, *Identification Plan for Land Mobile Stations*.

- b. To allow invalid GTP packets or packets that otherwise would fail parsing and be dropped, enter the following command:

```
hostname(config-pmap-p)# permit errors
```

By default, all invalid packets or packets that failed, during parsing, are dropped.

- c. To enable support for GSN pooling, use the **permit response** command.

If the security appliance performs GTP inspection, by default the security appliance drops GTP responses from GSNs that were not specified in the GTP request. This situation occurs when you use load-balancing among a pool of GSNs to provide efficiency and scalability of GPRS.

You can enable support for GSN pooling by using the **permit response** command. This command configures the security appliance to allow responses from any of a designated set of GSNs, regardless of the GSN to which a GTP request was sent. You identify the pool of load-balancing GSNs as a network object. Likewise, you identify the SGSN as a network object. If the GSN responding belongs to the same object group as the GSN that the GTP request was sent to and if the SGSN is in a object group that the responding GSN is permitted to send a GTP response to, the security appliance permits the response.

- d. To create an object to represent the pool of load-balancing GSNs, perform the following steps:

Use the **object-group** command to define a new network object group representing the pool of load-balancing GSNs.

```
hostname(config)# object-group network GSN-pool-name
hostname(config-network)#
```

For example, the following command creates an object group named gsnpool32:

```
hostname(config)# object-group network gsnpool32
hostname(config-network)#
```

- e. Use the **network-object** command to specify the load-balancing GSNs. You can do so with one **network-object** command per GSN, using the **host** keyword. You can also using **network-object** command to identify whole networks containing GSNs that perform load balancing.

```
hostname(config-network)# network-object host IP-address
```

For example, the following commands create three network objects representing individual hosts:

```
hostname(config-network)# network-object host 192.168.100.1
hostname(config-network)# network-object host 192.168.100.2
hostname(config-network)# network-object host 192.168.100.3
hostname(config-network)#
```

- f. To create an object to represent the SGSN that the load-balancing GSNs are permitted to respond to, perform the following steps:

- a. Use the **object-group** command to define a new network object group that will represent the SGSN that sends GTP requests to the GSN pool.

```
hostname(config)# object-group network SGSN-name
hostname(config-network)#
```

For example, the following command creates an object group named sgsn32:

```
hostname(config)# object-group network sgsn32
hostname(config-network)#
```

- b. Use the **network-object** command with the **host** keyword to identify the SGSN.

```
hostname(config-network)# network-object host IP-address
```

For example, the following command creates a network objects representing the SGSN:

```
hostname(config-network)# network-object host 192.168.50.100
hostname(config-network)#
```

- g. To allow GTP responses from any GSN in the network object representing the GSN pool, defined in c., d, to the network object representing the SGSN, defined in c., f., enter the following commands:

```
hostname(config)# gtp-map map_name
hostname(config-gtp-map)# permit response to-object-group SGSN-name from-object-group GSN-pool-name
```

For example, the following command permits GTP responses from any host in the object group named gsnpool32 to the host in the object group named sgsn32:

```
hostname(config-gtp-map)# permit response to-object-group sgsn32 from-object-group gsnpool32
```

The following example shows how to support GSN pooling by defining network objects for the GSN pool and the SGSN. An entire Class C network is defined as the GSN pool but you can identify multiple individual IP addresses, one per **network-object** command, instead of identifying whole networks. The example then modifies a GTP map to permit responses from the GSN pool to the SGSN.

```
hostname(config)# object-group network gsnpool32
hostname(config-network)# network-object 192.168.100.0 255.255.255.0
hostname(config)# object-group network sgsn32
hostname(config-network)# network-object host 192.168.50.100
```

```
hostname(config)# gtp-map gtp-policy
hostname(config-gtp-map)# permit response to-object-group sgsn32 from-object-group
gsnpool32
```

- h. To specify the maximum number of GTP requests that will be queued waiting for a response, enter the following command:

```
hostname(config-gtp-map)# request-queue max_requests
```

where the *max\_requests* argument sets the maximum number of GTP requests that will be queued waiting for a response, from 1 to 4294967295. The default is 200.

When the limit has been reached and a new request arrives, the request that has been in the queue for the longest time is removed. The Error Indication, the Version Not Supported and the SGSN Context Acknowledge messages are not considered as requests and do not enter the request queue to wait for a response.

- i. To change the inactivity timers for a GTP session, enter the following command:

```
hostname(config-gtp-map)# timeout {gsn | pdp-context | request | signaling | tunnel}
hh:mm:ss
```

Enter this command separately for each timeout.

The **gsn** keyword specifies the period of inactivity after which a GSN will be removed.

The **pdp-context** keyword specifies the maximum period of time allowed before beginning to receive the PDP context.

The **request** keyword specifies the maximum period of time allowed before beginning to receive the GTP message.

The **signaling** keyword specifies the period of inactivity after which the GTP signaling will be removed.

The **tunnel** keyword specifies the period of inactivity after which the GTP tunnel will be torn down.

The *hh:mm:ss* argument is the timeout where *hh* specifies the hour, *mm* specifies the minutes, and *ss* specifies the seconds. The value **0** means never tear down.

- j. To specify the maximum number of GTP tunnels allowed to be active on the security appliance, enter the following command:

```
hostname(config-gtp-map)# tunnel-limit max_tunnels
```

where the *max\_tunnels* argument is the maximum number of tunnels allowed, from 1 to 4294967295. The default is 500.

New requests will be dropped once the number of tunnels specified by this command is reached.

The following example shows how to limit the number of tunnels in the network:

```
hostname(config)# policy-map type inspect gtp gmap
hostname(config-pmap)# parameters
hostname(config-pmap-p)# tunnel-limit 3000

hostname(config)# policy-map global_policy
hostname(config-pmap)# class inspection_default
hostname(config-pmap-c)# inspect gtp gmap

hostname(config)# service-policy global_policy global
```

## Verifying and Monitoring GTP Inspection

To display GTP configuration, enter the **show service-policy inspect gtp** command in privileged EXEC mode. For the detailed syntax for this command, see the command page in the *Cisco Security Appliance Command Reference*.

Use the **show service-policy inspect gtp statistics** command to show the statistics for GTP inspection. The following is sample output from the **show service-policy inspect gtp statistics** command:

```
hostname# show service-policy inspect gtp statistics
GPRS GTP Statistics:
 version_not_support 0 msg_too_short 0
 unknown_msg 0 unexpected_sig_msg 0
 unexpected_data_msg 0 ie_duplicated 0
 mandatory_ie_missing 0 mandatory_ie_incorrect 0
 optional_ie_incorrect 0 ie_unknown 0
 ie_out_of_order 0 ie_unexpected 0
 total_forwarded 0 total_dropped 0
 signalling_msg_dropped 0 data_msg_dropped 0
 signalling_msg_forwarded 0 data_msg_forwarded 0
 total_created_pdp 0 total_deleted_pdp 0
 total_created_pdpmbc 0 total_deleted_pdpmbc 0
 pdp_non_existent 0
```

You can use the vertical bar (|) to filter the display. Type ?| for more display filtering options.

The following is sample GSN output from the **show service-policy inspect gtp statistics gsn** command:

```
hostname# show service-policy inspect gtp statistics gsn 9.9.9.9
1 in use, 1 most used, timeout 0:00:00

GTP GSN Statistics for 9.9.9.9, Idle 0:00:00, restart counter 0
 Tunnels Active 0Tunnels Created 0
 Tunnels Destroyed 0
 Total Messages Received 2
 Signaling Messages Data Messages
 total received 2 0
 dropped 0 0
 forwarded 2 0
```

Use the **show service-policy inspect gtp pdp-context** command to display PDP context-related information. The following is sample output from the **show service-policy inspect gtp pdp-context** command:

```
hostname# show service-policy inspect gtp pdp-context detail
1 in use, 1 most used, timeout 0:00:00

Version TID MS Addr SGSN Addr Idle APN
v1 1234567890123425 10.0.1.1 10.0.0.2 0:00:13 gprs.cisco.com

 user_name (IMSI): 214365870921435 MS address: 1.1.1.1
 primary pdp: Y
 sgsn_addr_signal: 10.0.0.2 sgsn_addr_data: 10.0.0.2
 ggsn_addr_signal: 10.1.1.1 ggsn_addr_data: 10.1.1.1
 sgsn control teid: 0x000001d1 sgsn data teid: 0x000001d3
 ggsn control teid: 0x6306ffa0 ggsn data teid: 0x6305f9fc
 seq_tpdu_up: 0 seq_tpdu_down: 0
 signal_sequence: 0
 upstream_signal_flow: 0 upstream_data_flow: 0
 downstream_signal_flow: 0 downstream_data_flow: 0
 RAupdate_flow: 0
```

The PDP context is identified by the tunnel ID, which is a combination of the values for IMSI and NSAPI. A GTP tunnel is defined by two associated PDP contexts in different GSN nodes and is identified with a Tunnel ID. A GTP tunnel is necessary to forward packets between an external packet data network and a MS user.

You can use the vertical bar (|) to filter the display, as in the following example:

```
hostname# show service-policy gtp statistics | grep gsn
```

## H.323 Inspection

This section describes the H.323 application inspection. This section includes the following topics:

- [H.323 Inspection Overview, page 25-38](#)
- [How H.323 Works, page 25-38](#)
- [Limitations and Restrictions, page 25-39](#)
- [Configuring H.323 and H.225 Timeout Values, page 25-42](#)
- [Verifying and Monitoring H.323 Inspection, page 25-43](#)

## H.323 Inspection Overview

H.323 inspection provides support for H.323 compliant applications such as Cisco CallManager and VocalTec Gatekeeper. H.323 is a suite of protocols defined by the International Telecommunication Union for multimedia conferences over LANs. The security appliance supports H.323 through Version 4, including H.323 v3 feature Multiple Calls on One Call Signaling Channel.

With H323 inspection enabled, the security appliance supports multiple calls on the same call signaling channel, a feature introduced with H.323 Version 3. This feature reduces call setup time and reduces the use of ports on the security appliance.

The two major functions of H.323 inspection are as follows:

- NAT the necessary embedded IPv4 addresses in the H.225 and H.245 messages. Because H.323 messages are encoded in PER encoding format, the security appliance uses an ASN.1 decoder to decode the H.323 messages.
- Dynamically allocate the negotiated H.245 and RTP/RTCP connections.

## How H.323 Works

The H.323 collection of protocols collectively may use up to two TCP connection and four to six UDP connections. FastConnect uses only one TCP connection, and RAS uses a single UDP connection for registration, admissions, and status.

An H.323 client may initially establish a TCP connection to an H.323 server using TCP port 1720 to request Q.931 call setup. As part of the call setup process, the H.323 terminal supplies a port number to the client to use for an H.245 TCP connection. In environments where H.323 gatekeeper is in use, the initial packet is transmitted using UDP.

H.323 inspection monitors the Q.931 TCP connection to determine the H.245 port number. If the H.323 terminals are not using FastConnect, the security appliance dynamically allocates the H.245 connection based on the inspection of the H.225 messages.



Within each H.245 message, the H.323 endpoints exchange port numbers that are used for subsequent UDP data streams. H.323 inspection inspects the H.245 messages to identify these ports and dynamically creates connections for the media exchange. RTP uses the negotiated port number, while RTCP uses the next higher port number.

The H.323 control channel handles H.225 and H.245 and H.323 RAS. H.323 inspection uses the following ports.

- 1718—Gate Keeper Discovery UDP port
- 1719—RAS UDP port
- 1720—TCP Control Port

You must permit traffic for the well-known H.323 port 1720 for the H.225 call signaling; however, the H.245 signaling ports are negotiated between the endpoints in the H.225 signaling. When an H.323 gatekeeper is used, the security appliance opens an H.225 connection based on inspection of the ACF message.

After inspecting the H.225 messages, the security appliance opens the H.245 channel and then inspects traffic sent over the H.245 channel as well. All H.245 messages passing through the security appliance undergo H.245 application inspection, which translates embedded IP addresses and opens the media channels negotiated in H.245 messages.

The H.323 ITU standard requires that a TPKT header, defining the length of the message, precede the H.225 and H.245, before being passed on to the reliable connection. Because the TPKT header does not necessarily need to be sent in the same TCP packet as H.225 and H.245 messages, the security appliance must remember the TPKT length to process and decode the messages properly. For each connection, the security appliance keeps a record that contains the TPKT length for the next expected message.

If the security appliance needs to perform NAT on IP addresses in messages, it changes the checksum, the UUIE length, and the TPKT, if it is included in the TCP packet with the H.225 message. If the TPKT is sent in a separate TCP packet, the security appliance proxy ACKs that TPKT and appends a new TPKT to the H.245 message with the new length.

**Note**

The security appliance does not support TCP options in the Proxy ACK for the TPKT.

Each UDP connection with a packet going through H.323 inspection is marked as an H.323 connection and times out with the H.323 timeout as configured with the **timeout** command.

## Limitations and Restrictions

The following are some of the known issues and limitations when using H.323 application inspection:

- Static PAT may not properly translate IP addresses embedded in optional fields within H.323 messages. If you experience this kind of problem, do not use static PAT with H.323.
- H.323 application inspection is not supported with NAT between same-security-level interfaces.
- When a NetMeeting client registers with an H.323 gatekeeper and tries to call an H.323 gateway that is also registered with the H.323 gatekeeper, the connection is established but no voice is heard in either direction. This problem is unrelated to the security appliance.
- If you configure a network static address where the network static address is the same as a third-party netmask and address, then any outbound H.323 connection fails.

## Configuring an H.323 Inspection Policy Map for Additional Inspection Control

To specify actions when a message violates a parameter, create an H.323 inspection policy map. You can then apply the inspection policy map when you enable H.323 inspection according to the [“Configuring Application Inspection” section on page 25-5](#).

To create an H.323 inspection policy map, perform the following steps:

- 
- Step 1** (Optional) Add one or more regular expressions for use in traffic matching commands according to the [“Creating a Regular Expression” section on page 21-12](#). See the types of text you can match in the **match** commands described in [Step 3](#).
- Step 2** (Optional) Create one or more regular expression class maps to group regular expressions according to the [“Creating a Regular Expression Class Map” section on page 21-14](#).
- Step 3** (Optional) Create an H.323 inspection class map by performing the following steps.

A class map groups multiple traffic matches. Traffic must match *all* of the **match** commands to match the class map. You can alternatively identify **match** commands directly in the policy map. The difference between creating a class map and defining the traffic match directly in the inspection policy map is that the class map lets you create more complex match criteria, and you can reuse class maps.

To specify traffic that should not match the class map, use the **match not** command. For example, if the **match not** command specifies the string “example.com,” then any traffic that includes “example.com” does not match the class map.

For the traffic that you identify in this class map, you can specify actions such as drop-connection, reset, and/or log the connection in the inspection policy map.

If you want to perform different actions for each **match** command, you should identify the traffic directly in the policy map.

- a. Create the class map by entering the following command:

```
hostname(config)# class-map type inspect h323 [match-all] class_map_name
hostname(config-cmap)#
```

Where *the class\_map\_name* is the name of the class map. The **match-all** keyword specifies that traffic must match all criteria to match the class map. **match-all** is the default and only option. The CLI enters class-map configuration mode, where you can enter one or more **match** commands.

- b. (Optional) To add a description to the class map, enter the following command:

```
hostname(config-cmap)# description string
```

Where *string* is the description of the class map (up to 200 characters).

- c. (Optional) To match a called party, enter the following command:

```
hostname(config-cmap)# match [not] called-party regex {class class_name | regex_name}
```

Where the **regex regex\_name** argument is the regular expression you created in [Step 1](#). The **class regex\_class\_name** is the regular expression class map you created in [Step 2](#).

- d. (Optional) To match a media type, enter the following command:

```
hostname(config-cmap)# match [not] media-type {audio | data | video}
```

- Step 4** Create an H.323 inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect h323 policy_map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

- Step 5** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

- Step 6** To apply actions to matching traffic, perform the following steps.

- a. Specify the traffic on which you want to perform actions using one of the following methods:

- Specify the H.323 class map that you created in [Step 3](#) by entering the following command:

```
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

- Specify traffic directly in the policy map using one of the **match** commands described in [Step 3](#). If you use a **match not** command, then any traffic that does not match the criterion in the **match not** command has the action applied.

- b. Specify the action you want to perform on the matching traffic by entering the following command:

```
hostname(config-pmap-c)# {[drop [send-protocol-error] |
drop-connection [send-protocol-error] | mask | reset] [log] | rate-limit message_rate}
```

Not all options are available for each **match** or **class** command. See the CLI help or the *Cisco Security Appliance Command Reference* for the exact options available.

The **drop** keyword drops all packets that match.

The **send-protocol-error** keyword sends a protocol error message.

The **drop-connection** keyword drops the packet and closes the connection.

The **mask** keyword masks out the matching portion of the packet.

The **reset** keyword drops the packet, closes the connection, and sends a TCP reset to the server and/or client.

The **log** keyword, which you can use alone or with one of the other keywords, sends a system log message.

The **rate-limit** *message\_rate* argument limits the rate of messages.

You can specify multiple **class** or **match** commands in the policy map. For information about the order of **class** and **match** commands, see the “[Defining Actions in an Inspection Policy Map](#)” section on [page 21-8](#).

- Step 7** To configure parameters that affect the inspection engine, perform the following steps:

- a. To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap)# parameters
hostname(config-pmap-p)#
```

- b. To define the H.323 call duration limit, enter the following command:

```
hostname(config-pmap-p)# call-duration-limit time
```

Where *time* is the call duration limit in seconds. Range is from 0:0:0 to 1163:0:0. A value of 0 means never timeout.

- c. To enforce call party number used in call setup, enter the following command:

```
hostname(config-pmap-p)# call-party-number
```

- d. To enforce H.245 tunnel blocking, enter the following command:

```
hostname(config-pmap-p)# h245-tunnel-block action {drop-connection | log}
```

- e. To define an hsi group and enter hsi group configuration mode, enter the following command:

```
hostname(config-pmap-p)# hsi-group id
```

Where *id* is the hsi group ID. Range is from 0 to 2147483647.

To add an hsi to the hsi group, enter the following command in hsi group configuration mode:

```
hostname(config-h225-map-hsi-grp)# hsi ip_address
```

Where *ip\_address* is the host to add. A maximum of five hosts per hsi group are allowed.

To add an endpoint to the hsi group, enter the following command in hsi group configuration mode:

```
hostname(config-h225-map-hsi-grp)# endpoint ip_address if_name
```

Where *ip\_address* is the endpoint to add and *if\_name* is the interface through which the endpoint is connected to the security appliance. A maximum of ten endpoints per hsi group are allowed.

- f. To check RTP packets flowing on the pinholes for protocol conformance, enter the following command:

```
hostname(config-pmap-p)# rtp-conformance [enforce-payloadtype]
```

Where the **enforce-payloadtype** keyword enforces the payload type to be audio or video based on the signaling exchange.

- g. To enable state checking validation, enter the following command:

```
hostname(config-pmap-p)# state-checking {h225 | ras}
```

The following example shows how to configure phone number filtering:

```
hostname(config)# regex caller 1 "5551234567"
hostname(config)# regex caller 2 "5552345678"
hostname(config)# regex caller 3 "5553456789"

hostname(config)# class-map type inspect h323 match-all h323_traffic
hostname(config-pmap-c)# match called-party regex caller1
hostname(config-pmap-c)# match calling-party regex caller2

hostname(config)# policy-map type inspect h323 h323_map
hostname(config-pmap)# parameters
hostname(config-pmap-p)# class h323_traffic
hostname(config-pmap-c)# drop
```

## Configuring H.323 and H.225 Timeout Values

To configure the idle time after which an H.225 signalling connection is closed, use the **timeout h225** command. The default for H.225 timeout is one hour.

To configure the idle time after which an H.323 control connection is closed, use the **timeout h323** command. The default is five minutes.

## Verifying and Monitoring H.323 Inspection

This section describes how to display information about H.323 sessions. This section includes the following topics:

- [Monitoring H.225 Sessions, page 25-43](#)
- [Monitoring H.245 Sessions, page 25-43](#)
- [Monitoring H.323 RAS Sessions, page 25-44](#)

### Monitoring H.225 Sessions

The **show h225** command displays information for H.225 sessions established across the security appliance. Along with the **debug h323 h225 event**, **debug h323 h245 event**, and **show local-host** commands, this command is used for troubleshooting H.323 inspection engine issues.

Before entering the **show h225**, **show h245**, or **show h323-ras** commands, we recommend that you configure the **pager** command. If there are a lot of session records and the **pager** command is not configured, it may take a while for the **show** command output to reach its end. If there is an abnormally large number of connections, check that the sessions are timing out based on the default timeout values or the values set by you. If they are not, then there is a problem that needs to be investigated.

The following is sample output from the **show h225** command:

```
hostname# show h225
Total H.323 Calls: 1
1 Concurrent Call(s) for
 Local: 10.130.56.3/1040 Foreign: 172.30.254.203/1720
 1. CRV 9861
 Local: 10.130.56.3/1040 Foreign: 172.30.254.203/1720
0 Concurrent Call(s) for
 Local: 10.130.56.4/1050 Foreign: 172.30.254.205/1720
```

This output indicates that there is currently 1 active H.323 call going through the security appliance between the local endpoint 10.130.56.3 and foreign host 172.30.254.203, and for these particular endpoints, there is 1 concurrent call between them, with a CRV for that call of 9861.

For the local endpoint 10.130.56.4 and foreign host 172.30.254.205, there are 0 concurrent calls. This means that there is no active call between the endpoints even though the H.225 session still exists. This could happen if, at the time of the **show h225** command, the call has already ended but the H.225 session has not yet been deleted. Alternately, it could mean that the two endpoints still have a TCP connection opened between them because they set “maintainConnection” to TRUE, so the session is kept open until they set it to FALSE again, or until the session times out based on the H.225 timeout value in your configuration.

### Monitoring H.245 Sessions

The **show h245** command displays information for H.245 sessions established across the security appliance by endpoints using slow start. Slow start is when the two endpoints of a call open another TCP control channel for H.245. Fast start is where the H.245 messages are exchanged as part of the H.225 messages on the H.225 control channel.) Along with the **debug h323 h245 event**, **debug h323 h225 event**, and **show local-host** commands, this command is used for troubleshooting H.323 inspection engine issues.

The following is sample output from the **show h245** command:

```
hostname# show h245
Total: 1
```

```

1 LOCAL TPKT FOREIGN TPKT
 10.130.56.3/1041 0 172.30.254.203/1245 0
 MEDIA: LCN 258 Foreign 172.30.254.203 RTP 49608 RTCP 49609
 Local 10.130.56.3 RTP 49608 RTCP 49609
 MEDIA: LCN 259 Foreign 172.30.254.203 RTP 49606 RTCP 49607
 Local 10.130.56.3 RTP 49606 RTCP 49607

```

There is currently one H.245 control session active across the security appliance. The local endpoint is 10.130.56.3, and we are expecting the next packet from this endpoint to have a TPKT header because the TPKT value is 0. The TKTP header is a 4-byte header preceding each H.225/H.245 message. It gives the length of the message, including the 4-byte header. The foreign host endpoint is 172.30.254.203, and we are expecting the next packet from this endpoint to have a TPKT header because the TPKT value is 0.

The media negotiated between these endpoints have an LCN of 258 with the foreign RTP IP address/port pair of 172.30.254.203/49608 and an RTCP IP address/port of 172.30.254.203/49609 with a local RTP IP address/port pair of 10.130.56.3/49608 and an RTCP port of 49609.

The second LCN of 259 has a foreign RTP IP address/port pair of 172.30.254.203/49606 and an RTCP IP address/port pair of 172.30.254.203/49607 with a local RTP IP address/port pair of 10.130.56.3/49606 and RTCP port of 49607.

## Monitoring H.323 RAS Sessions

The **show h323-ras** command displays information for H.323 RAS sessions established across the security appliance between a gatekeeper and its H.323 endpoint. Along with the **debug h323 ras event** and **show local-host** commands, this command is used for troubleshooting H.323 RAS inspection engine issues.

The **show h323-ras** command displays connection information for troubleshooting H.323 inspection engine issues. The following is sample output from the **show h323-ras** command:

```

hostname# show h323-ras
Total: 1
 GK Caller
 172.30.254.214 10.130.56.14

```

This output shows that there is one active registration between the gatekeeper 172.30.254.214 and its client 10.130.56.14.

# HTTP Inspection

This section describes the HTTP inspection engine. This section includes the following topics:

- [HTTP Inspection Overview, page 25-44](#)
- [Configuring an HTTP Inspection Policy Map for Additional Inspection Control, page 25-45](#)

## HTTP Inspection Overview

Use the HTTP inspection engine to protect against specific attacks and other threats that may be associated with HTTP traffic. HTTP inspection performs several functions:

- Enhanced HTTP inspection
- URL screening through N2H2 or Websense
- Java and ActiveX filtering

The latter two features are configured in conjunction with the **filter** command. For more information about filtering, see [Chapter 20, “Applying Filtering Services.”](#)

**Note**

The **no inspect http** command also disables the **filter url** command.

The enhanced HTTP inspection feature, which is also known as an application firewall and is available when you configure an HTTP map (see [“Configuring an HTTP Inspection Policy Map for Additional Inspection Control”](#)), can help prevent attackers from using HTTP messages for circumventing network security policy. It verifies the following for all HTTP messages:

- Conformance to RFC 2616
- Use of RFC-defined methods only.
- Compliance with the additional criteria.

## Configuring an HTTP Inspection Policy Map for Additional Inspection Control

To specify actions when a message violates a parameter, create an HTTP inspection policy map. You can then apply the inspection policy map when you enable HTTP inspection according to the [“Configuring Application Inspection”](#) section on page 25-5.

**Note**

When you enable HTTP inspection with an inspection policy map, strict HTTP inspection with the action reset and log is enabled by default. You can change the actions performed in response to inspection failure, but you cannot disable strict inspection as long as the inspection policy map remains enabled.

To create an HTTP inspection policy map, perform the following steps:

- Step 1** (Optional) Add one or more regular expressions for use in traffic matching commands according to the [“Creating a Regular Expression”](#) section on page 21-12. See the types of text you can match in the **match** commands described in [Step 3](#).
- Step 2** (Optional) Create one or more regular expression class maps to group regular expressions according to the [“Creating a Regular Expression Class Map”](#) section on page 21-14.
- Step 3** (Optional) Create an HTTP inspection class map by performing the following steps.

A class map groups multiple traffic matches. Traffic must match *all* of the **match** commands to match the class map. You can alternatively identify **match** commands directly in the policy map. The difference between creating a class map and defining the traffic match directly in the inspection policy map is that the class map lets you create more complex match criteria, and you can reuse class maps.

To specify traffic that should not match the class map, use the **match not** command. For example, if the **match not** command specifies the string “example.com,” then any traffic that includes “example.com” does not match the class map.

For the traffic that you identify in this class map, you can specify actions such as drop, drop-connection, reset, mask, set the rate limit, and/or log the connection in the inspection policy map.

If you want to perform different actions for each **match** command, you should identify the traffic directly in the policy map.

- a. Create the class map by entering the following command:

```
hostname(config)# class-map type inspect http [match-all] class_map_name
hostname(config-cmap)#
```

Where *class\_map\_name* is the name of the class map. The **match-all** keyword specifies that traffic must match all criteria to match the class map. **match-all** is the default and only option. The CLI enters class-map configuration mode, where you can enter one or more **match** commands.

- b. (Optional) To add a description to the class map, enter the following command:

```
hostname(config-cmap)# description string
```

- c. (Optional) To match traffic with a content-type field in the HTTP response that does not match the accept field in the corresponding HTTP request message, enter the following command:

```
hostname(config-cmap)# match [not] req-resp content-type mismatch
```

- d. (Optional) To match text found in the HTTP request message arguments, enter the following command:

```
hostname(config-cmap)# match [not] request args regex [regex_name | class
regex_class_name]
```

Where the *regex\_name* is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#).

- e. (Optional) To match text found in the HTTP request message body or to match traffic that exceeds the maximum HTTP request message body length, enter the following command:

```
hostname(config-cmap)# match [not] request body {regex [regex_name | class
regex_class_name] | length gt max_bytes}
```

Where the **regex** *regex\_name* argument is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#). The **length** **gt** *max\_bytes* is the maximum message body length in bytes.

- f. (Optional) To match text found in the HTTP request message header, or to restrict the count or length of the header, enter the following command:

```
hostname(config-cmap)# match [not] request header {[field]
[regex [regex_name | class regex_class_name]] |
[length gt max_length_bytes | count gt max_count_bytes]}
```

Where the *field* is the predefined message header keyword. The **regex** *regex\_name* argument is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#). The **length** **gt** *max\_bytes* is the maximum message body length in bytes. The **count** **gt** *max\_count* is the maximum number of header fields.

- g. (Optional) To match text found in the HTTP request message method, enter the following command:

```
hostname(config-cmap)# match [not] request method {[method] |
[regex [regex_name | class regex_class_name]]}
```

Where the *method* is the predefined message method keyword. The **regex** *regex\_name* argument is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#).

- h. (Optional) To match text found in the HTTP request message URI, enter the following command:

```
hostname(config-cmap)# match [not] request uri {regex [regex_name | class
regex_class_name] | length gt max_bytes}
```

Where the **regex** *regex\_name* argument is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#). The **length** **gt** *max\_bytes* is the maximum message body length in bytes.



- i. Optional) To match text found in the HTTP response message body, or to comment out Java applet and Active X object tags in order to filter them, enter the following command:

```
hostname(config-cmap)# match [not] response body {[active-x] | [java-applet] |
[regex [regex_name] | class regex_class_name]} | length gt max_bytes}
```

Where the **regex** *regex\_name* argument is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#). The **length gt** *max\_bytes* is the maximum message body length in bytes.

- j. (Optional) To match text found in the HTTP response message header, or to restrict the count or length of the header, enter the following command:

```
hostname(config-cmap)# match [not] response header {[field]
[regex [regex_name] | class regex_class_name]} |
[length gt max_length_bytes | count gt max_count]}
```

Where the *field* is the predefined message header keyword. The **regex** *regex\_name* argument is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#). The **length gt** *max\_bytes* is the maximum message body length in bytes. The **count gt** *max\_count* is the maximum number of header fields.

- k. (Optional) To match text found in the HTTP response message status line, enter the following command:

```
hostname(config-cmap)# match [not] response status-line {regex [regex_name] | class
regex_class_name}
```

Where the **regex** *regex\_name* argument is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#).

- Step 4** Create an HTTP inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect http policy_map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

- Step 5** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

- Step 6** To apply actions to matching traffic, perform the following steps.

- a. Specify the traffic on which you want to perform actions using one of the following methods:

- Specify the HTTP class map that you created in [Step 3](#) by entering the following command:

```
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

- Specify traffic directly in the policy map using one of the **match** commands described in [Step 3](#). If you use a **match not** command, then any traffic that does not match the criterion in the **match not** command has the action applied.

- b. Specify the action you want to perform on the matching traffic by entering the following command:

```
hostname(config-pmap-c)# {[drop [send-protocol-error] |
drop-connection [send-protocol-error] | mask | reset] [log] | rate-limit message_rate}
```

Not all options are available for each **match** or **class** command. See the CLI help or the *Cisco Security Appliance Command Reference* for the exact options available.

The **drop** keyword drops all packets that match.

The **send-protocol-error** keyword sends a protocol error message.

The **drop-connection** keyword drops the packet and closes the connection.

The **mask** keyword masks out the matching portion of the packet.

The **reset** keyword drops the packet, closes the connection, and sends a TCP reset to the server and/or client.

The **log** keyword, which you can use alone or with one of the other keywords, sends a system log message.

The **rate-limit** *message\_rate* argument limits the rate of messages.

You can specify multiple **class** or **match** commands in the policy map. For information about the order of **class** and **match** commands, see the “[Defining Actions in an Inspection Policy Map](#)” section on [page 21-8](#).

**Step 7** To configure parameters that affect the inspection engine, perform the following steps:

- a. To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap)# parameters
hostname(config-pmap-p)#
```

- b. To check for HTTP protocol violations, enter the following command:

```
hostname(config-pmap-p)# protocol-violation [action [drop-connection / reset / log]]
```

Where the **drop-connection** action closes the connection. The **reset** action closes the connection and sends a TCP reset to the client. The **log** action sends a system log message when this policy map matches traffic.

- c. To substitute a string for the server header field, enter the following command:

```
hostname(config-pmap-p)# spoof-server string
```

Where the *string* argument is the string to substitute for the server header field. Note: WebVPN streams are not subject to the **spoof-server** command.

The following example shows how to define an HTTP inspection policy map that will allow and log any HTTP connection that attempts to access “www.xyz.com/\*.asp” or “www.xyz[0-9][0-9].com” with methods “GET” or “PUT.” All other URL/Method combinations will be silently allowed.

```
hostname(config)# regex url1 "www\.xyz\.com/.*\.asp"
hostname(config)# regex url2 "www\.xyz[0-9][0-9]\.com"
hostname(config)# regex get "GET"
hostname(config)# regex put "PUT"

hostname(config)# class-map type regex match-any url_to_log
hostname(config-cmap)# match regex url1
hostname(config-cmap)# match regex url2
hostname(config-cmap)# exit

hostname(config)# class-map type regex match-any methods_to_log
hostname(config-cmap)# match regex get
hostname(config-cmap)# match regex put
hostname(config-cmap)# exit

hostname(config)# class-map type inspect http http_url_policy
hostname(config-cmap)# match request uri regex class url_to_log
hostname(config-cmap)# match request method regex class methods_to_log
hostname(config-cmap)# exit
```

```
hostname(config)# policy-map type inspect http http_policy
hostname(config-pmap)# class http_url_policy
hostname(config-pmap-c)# log
```

## Instant Messaging Inspection

This section describes the IM inspection engine. This section includes the following topics:

- [IM Inspection Overview, page 25-49](#)
- [Configuring an Instant Messaging Inspection Policy Map for Additional Inspection Control, page 25-49](#)

### IM Inspection Overview

The IM inspect engine lets you apply fine grained controls on the IM application to control the network usage and stop leakage of confidential data, propagation of worms, and other threats to the corporate network.

### Configuring an Instant Messaging Inspection Policy Map for Additional Inspection Control

To specify actions when a message violates a parameter, create an IM inspection policy map. You can then apply the inspection policy map when you enable IM inspection according to the [“Configuring Application Inspection” section on page 25-5](#).

To create an IM inspection policy map, perform the following steps:

- 
- Step 1** (Optional) Add one or more regular expressions for use in traffic matching commands according to the [“Creating a Regular Expression” section on page 21-12](#). See the types of text you can match in the **match** commands described in [Step 3](#).
- Step 2** (Optional) Create one or more regular expression class maps to group regular expressions according to the [“Creating a Regular Expression Class Map” section on page 21-14](#).
- Step 3** (Optional) Create an IM inspection class map by performing the following steps.

A class map groups multiple traffic matches. Traffic must match *all* of the **match** commands to match the class map. You can alternatively identify **match** commands directly in the policy map. The difference between creating a class map and defining the traffic match directly in the inspection policy map is that the class map lets you create more complex match criteria, and you can reuse class maps.

To specify traffic that should not match the class map, use the **match not** command. For example, if the **match not** command specifies the string “example.com,” then any traffic that includes “example.com” does not match the class map.

For the traffic that you identify in this class map, you can specify actions such as drop-connection, reset, and/or log the connection in the inspection policy map.

If you want to perform different actions for each **match** command, you should identify the traffic directly in the policy map.

- a. Create the class map by entering the following command:

```
hostname(config)# class-map type inspect im [match-all] class_map_name
hostname(config-cmap)#
```

Where *the class\_map\_name* is the name of the class map. The **match-all** keyword specifies that traffic must match all criteria to match the class map. **match-all** is the default and only option. The CLI enters class-map configuration mode, where you can enter one or more **match** commands.

- b. (Optional) To add a description to the class map, enter the following command:

```
hostname(config-cmap)# description string
```

Where *the string* is the description of the class map (up to 200 characters).

- c. (Optional) To match traffic of a specific IM protocol, such as Yahoo or MSN, enter the following command:

```
hostname(config-cmap)# match [not] protocol {im-yahoo | im-msn}
```

- d. (Optional) To match a specific IM service, such as chat, file-transfer, webcam, voice-chat, conference, or games, enter the following command:

```
hostname(config-cmap)# match [not] service {chat | file-transfer | webcam | voice-chat | conference | games}
```

- e. (Optional) To match the source login name of the IM message, enter the following command:

```
hostname(config-cmap)# match [not] login-name regex {class class_name | regex_name}
```

Where the **regex regex\_name** argument is the regular expression you created in [Step 1](#). The **class regex\_class\_name** is the regular expression class map you created in [Step 2](#).

- f. (Optional) To match the destination login name of the IM message, enter the following command:

```
hostname(config-cmap)# match [not] peer-login-name regex {class class_name | regex_name}
```

Where the **regex regex\_name** argument is the regular expression you created in [Step 1](#). The **class regex\_class\_name** is the regular expression class map you created in [Step 2](#).

- g. (Optional) To match the source IP address of the IM message, enter the following command:

```
hostname(config-cmap)# match [not] ip-address ip_address ip_address_mask
```

Where the *ip\_address* and the *ip\_address\_mask* is the IP address and netmask of the message source.

- h. (Optional) To match the destination IP address of the IM message, enter the following command:

```
hostname(config-cmap)# match [not] peer-ip-address ip_address ip_address_mask
```

Where the *ip\_address* and the *ip\_address\_mask* is the IP address and netmask of the message destination.

- i. (Optional) To match the version of the IM message, enter the following command:

```
hostname(config-cmap)# match [not] version regex {class class_name | regex_name}
```

Where the **regex regex\_name** argument is the regular expression you created in [Step 1](#). The **class regex\_class\_name** is the regular expression class map you created in [Step 2](#).

- j. (Optional) To match the filename of the IM message, enter the following command:

```
hostname(config-cmap)# match [not] filename regex {class class_name | regex_name}
```

Where the **regex regex\_name** argument is the regular expression you created in [Step 1](#). The **class regex\_class\_name** is the regular expression class map you created in [Step 2](#).

**Note**

Not supported using MSN IM protocol.

**Step 4** Create an IM inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect im policy_map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

**Step 5** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

**Step 6** Specify the traffic on which you want to perform actions using one of the following methods:

- Specify the IM class map that you created in [Step 3](#) by entering the following command:

```
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

- Specify traffic directly in the policy map using one of the **match** commands described in [Step 3](#). If you use a **match not** command, then any traffic that does not match the criterion in the **match not** command has the action applied.

You can specify multiple **class** or **match** commands in the policy map. For information about the order of **class** and **match** commands, see the “[Defining Actions in an Inspection Policy Map](#)” section on [page 21-8](#).

**Step 7** Specify the action you want to perform on the matching traffic by entering the following command:

```
hostname(config-pmap-c)# {drop-connection | reset | log}
```

Where the **drop-connection** action closes the connection. The **reset** action closes the connection and sends a TCP reset to the client. The **log** action sends a system log message when this policy map matches traffic.

The following example shows how to define an IM inspection policy map.

```
hostname(config)# regex loginname1 "ying@yahoo.com"
hostname(config)# regex loginname2 "Kevin@yahoo.com"
hostname(config)# regex loginname3 "rahul@yahoo.com"
hostname(config)# regex loginname3 "darshant@yahoo.com"
hostname(config)# regex yhoos_version_regex "1\\.0"

hostname(config)# class-map type regex match-any yahoo_src_login_name_regex
hostname(config-cmap)# match regex loginname1
hostname(config-cmap)# match regex loginname2

hostname(config)# class-map type regex match-any yahoo_dst_login_name_regex
hostname(config-cmap)# match regex loginname3
hostname(config-cmap)# match regex loginname4

hostname(config)# class-map type regex match-any yhoos_file_block_list
hostname(config-cmap)# match regex ".*\\.gif"
hostname(config-cmap)# match regex ".*\\.exe"

hostname(config)# class-map type regex match-any new_im_regexp
hostname(config-cmap)# match regexp "new_im_regexp"
```

```

hostname(config)# class-map type inspect im match-all yahoo_im_policy
hostname(config-cmap)# match login-name regex class yhoosrc_login_name_regex
hostname(config-cmap)# match peer-login-name regex class yhoosrc_login_name_regex

hostname(config)# class-map type inspect im yahoo_im_policy2
hostname(config-cmap)# match version regex yahoo_version_regex

hostname(config)# class-map im_inspect_class_map
hostname(config-cmap)# match default-inspection-traffic

hostname(config)# policy-map type im im_policy_all
hostname(config-pmap)# class yahoo_in_file_xfer_policy
hostname(config-pmap-c)# drop-connection
hostname(config-pmap-c)# class yhoosrc_im_policy
hostname(config-pmap-c)# drop-connection
hostname(config-pmap-c)# class yhoosrc_im_policy2
hostname(config-pmap-c)# reset
hostname(config-pmap-c)# match im-pattern regex class new_im_regexp
hostname(config-pmap-c)# action log
hostname(config)# policy-map global_policy_name
hostname(config-pmap)# class im_inspection_class_map
hostname(config-pmap-c)# inspect im im_policy_all

```

## ICMP Inspection

The ICMP inspection engine allows ICMP traffic to have a “session” so it can be inspected like TCP and UDP traffic. Without the ICMP inspection engine, we recommend that you do not allow ICMP through the security appliance in an access list. Without stateful inspection, ICMP can be used to attack your network. The ICMP inspection engine ensures that there is only one response for each request, and that the sequence number is correct.

## ICMP Error Inspection

When this feature is enabled, the security appliance creates translation sessions for intermediate hops that send ICMP error messages, based on the NAT configuration. The security appliance overwrites the packet with the translated IP addresses.

When disabled, the security appliance does not create translation sessions for intermediate nodes that generate ICMP error messages. ICMP error messages generated by the intermediate nodes between the inside host and the security appliance reach the outside host without consuming any additional NAT resource. This is undesirable when an outside host uses the traceroute command to trace the hops to the destination on the inside of the security appliance. When the security appliance does not translate the intermediate hops, all the intermediate hops appear with the mapped destination IP address.

The ICMP payload is scanned to retrieve the five-tuple from the original packet. Using the retrieved five-tuple, a lookup is performed to determine the original address of the client. The ICMP error inspection engine makes the following changes to the ICMP packet:

- In the IP Header, the mapped IP is changed to the real IP (Destination Address) and the IP checksum is modified.
- In the ICMP Header, the ICMP checksum is modified due to the changes in the ICMP packet.
- In the Payload, the following changes are made:
  - Original packet mapped IP is changed to the real IP

- Original packet mapped port is changed to the real Port
- Original packet IP checksum is recalculated

## ILS Inspection

The ILS inspection engine provides NAT support for Microsoft NetMeeting, SiteServer, and Active Directory products that use LDAP to exchange directory information with an ILS server.

The security appliance supports NAT for ILS, which is used to register and locate endpoints in the ILS or SiteServer Directory. PAT cannot be supported because only IP addresses are stored by an LDAP database.

For search responses, when the LDAP server is located outside, NAT should be considered to allow internal peers to communicate locally while registered to external LDAP servers. For such search responses, xlates are searched first, and then DNAT entries to obtain the correct address. If both of these searches fail, then the address is not changed. For sites using NAT 0 (no NAT) and not expecting DNAT interaction, we recommend that the inspection engine be turned off to provide better performance.

Additional configuration may be necessary when the ILS server is located inside the security appliance border. This would require a hole for outside clients to access the LDAP server on the specified port, typically TCP 389.

Because ILS traffic only occurs on the secondary UDP channel, the TCP connection is disconnected after the TCP inactivity interval. By default, this interval is 60 minutes and can be adjusted using the **timeout** command.

ILS/LDAP follows a client/server model with sessions handled over a single TCP connection. Depending on the client's actions, several of these sessions may be created.

During connection negotiation time, a BIND PDU is sent from the client to the server. Once a successful BIND RESPONSE from the server is received, other operational messages may be exchanged (such as ADD, DEL, SEARCH, or MODIFY) to perform operations on the ILS Directory. The ADD REQUEST and SEARCH RESPONSE PDUs may contain IP addresses of NetMeeting peers, used by H.323 (SETUP and CONNECT messages) to establish the NetMeeting sessions. Microsoft NetMeeting v2.X and v3.X provides ILS support.

The ILS inspection performs the following operations:

- Decodes the LDAP REQUEST/RESPONSE PDUs using the BER decode functions
- Parses the LDAP packet
- Extracts IP addresses
- Translates IP addresses as necessary
- Encodes the PDU with translated addresses using BER encode functions
- Copies the newly encoded PDU back to the TCP packet
- Performs incremental TCP checksum and sequence number adjustment

ILS inspection has the following limitations:

- Referral requests and responses are not supported
- Users in multiple directories are not unified
- Single users having multiple identities in multiple directories cannot be recognized by NAT

**Note**

Because H225 call signalling traffic only occurs on the secondary UDP channel, the TCP connection is disconnected after the interval specified by the TCP **timeout** command. By default, this interval is set at 60 minutes.

## IPSec Pass Through Inspection

This section describes IPSec Pass Through application inspection. This section includes the following topics:

- [IPSec Pass Through Inspection Overview, page 25-54](#)
- [Configuring an IPSec Pass Through Inspection Policy Map for Additional Inspection Control, page 25-54](#)

## IPSec Pass Through Inspection Overview

The IPSec Pass Through inspection engine lets the security appliance pass ESP (IP protocol 50) and AH (IP protocol 51) traffic that is formed between two hosts because of successful IKE (UDP port 500) negotiation without the requirement of specific ESP or AH access lists.

The inspection engine works on IKE UDP port 500 to create the control flow. The **inspect ipsec-pass-thru** command is attached to an UDP flow as defined in the MPF framework. When an ESP or AH packet between the two peers arrives at the device, or an UDP packet with either source or destination port equal to 500, the packet is sent to the inspect module.

The ESP or AH traffic is permitted by the inspection engine with the configured idle timeout if there is an existing control flow and it is within the connection limit defined in the MPF framework. A new control flow is created for IKE UDP port 500 traffic with the configured UDP idle timeout if there isn't one, or it uses the existing flow.

To ensure that the packet arrives into the inspection engine, a hole is punched for all such traffic (ESP and AH). This inspect is attached to the control flow. The control flow is present as long as there is at least one data flow (ESP or AH) established, but the traffic always flows on the same connection. Since this IKE connection is kept open as long as data flows, a rekey would always succeed. The flows are created irrespective of NAT or no NAT.

**Note**

PAT is not supported.

## Configuring an IPSec Pass Through Inspection Policy Map for Additional Inspection Control

Inspect IPSec Pass Through is disabled by default. When enabled without using a parameter map, the inspection uses the default IPSec Pass Through parameter map, which allows only ESP traffic with unlimited connections and the default idle timeout of 10 minutes for the ESP connection.

To pass ESP or AH traffic, IPSec Pass Through parameter map is required.

To create an IPSec Pass Through map, perform the following steps:



**Step 1** To create an IPSec Pass Through inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect IPSec-pass-thru map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

**Step 2** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

**Step 3** To configure parameters that affect the inspection engine, perform the following steps:

a. To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap)# parameters
hostname(config-pmap-p)#
```

b. To configure for ESP traffic, enter the following command:

```
hostname(config-pmap-p)# esp per-client-max value timeout timeout
```

c. To configure for AH traffic, enter the following command:

```
hostname(config-pmap-p)# ah per-client-max value timeout timeout
```

The following example shows how to define an IPSec Pass Through map:

```
hostname(config)# access-list test-udp-acl extended permit udp any any eq 500

hostname(config)# class-map test-udp-class
hostname(config-cmap)# match access-list test-udp-acl

hostname(config)# policy-map type inspect IPSec-pass-thru IPsec-map
hostname(config-pmap)# parameters
hostname(config-pmap-p)# esp per-client-max 32 timeout 00:06:00
hostname(config-pmap-p)# ah per-client-max 16 timeout 00:05:00

hostname(config)# policy-map test-udp-policy
hostname(config-pmap)# class test-udp-class
hostname(config-pmap-c)# inspect IPSec-pass-thru IPsec-map
```

This policy is applied on the interface that has the policy to permit UDP 500 traffic through initially. In this example it is the outside interface.

To verify that the inspection engine opens the ESP or AH data flows for IPSec Pass Through based on the IKE control flow, use the **show conn** command:

```
hostname(config)# show conn
ESP out 192.168.51.25 in 192.168.52.49 idle 0:00:00 bytes 108
AH out 192.168.51.25 in 192.168.52.49 idle 0:00:01 bytes 0
ESP out 192.168.51.25 in 192.168.52.49 idle 0:00:01 bytes 0
UDP out 192.168.51.25:500 in 192.168.52.49:500 idle 0:00:00 flags -
AH out 192.168.51.25 in 192.168.52.50 idle 0:00:22 bytes 0
ESP out 192.168.51.25 in 192.168.52.50 idle 0:00:22 bytes 0
UDP out 192.168.51.25:500 in 192.168.52.50:500 idle 0:00:00 flags -
ESP out 192.168.51.25 in 192.168.52.49 idle 0:00:00 bytes 108
AH out 0.0.0.0 in 0.0.0.0 idle 0:00:01 bytes 0
ESP out 0.0.0.0 in 0.0.0.0 idle 0:00:01 bytes 0
AH out 192.168.51.25 in 192.168.52.50 idle 0:00:00 bytes 2080
AH out 0.0.0.0 in 0.0.0.0 idle 0:00:22 bytes 0
ESP out 0.0.0.0 in 0.0.0.0 idle 0:00:22 bytes 0
```

# MGCP Inspection

This section describes MGCP application inspection. This section includes the following topics:

- [MGCP Inspection Overview, page 25-56](#)
- [Configuring an MGCP Inspection Policy Map for Additional Inspection Control, page 25-58](#)
- [Configuring MGCP Timeout Values, page 25-59](#)
- [Verifying and Monitoring MGCP Inspection, page 25-59](#)

## MGCP Inspection Overview


MGCP is a master/slave protocol used to control media gateways from external call control elements called media gateway controllers or call agents. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Using NAT and PAT with MGCP lets you support a large number of devices on an internal network with a limited set of external (global) addresses. Examples of media gateways are:

- Trunking gateways, that interface between the telephone network and a Voice over IP network. Such gateways typically manage a large number of digital circuits.
- Residential gateways, that provide a traditional analog (RJ11) interface to a Voice over IP network. Examples of residential gateways include cable modem/cable set-top boxes, xDSL devices, broad-band wireless devices.
- Business gateways, that provide a traditional digital PBX interface or an integrated soft PBX interface to a Voice over IP network.

**Note**

To avoid policy failure when upgrading from ASA version 7.1, all layer 7 and layer 3 policies must have distinct names. For instance, a previously configured policy map with the same name as a previously configured MGCP map must be changed before the upgrade.

MGCP messages are transmitted over UDP. A response is sent back to the source address (IP address and UDP port number) of the command, but the response may not arrive from the same address as the command was sent to. This can happen when multiple call agents are being used in a failover configuration and the call agent that received the command has passed control to a backup call agent, which then sends the response. [Figure 25-4](#) illustrates how NAT can be used with MGCP.

**Figure 25-4 Using NAT with MGCP**

MGCP endpoints are physical or virtual sources and destinations for data. Media gateways contain endpoints on which the call agent can create, modify and delete connections to establish and control media sessions with other multimedia endpoints. Also, the call agent can instruct the endpoints to detect certain events and generate signals. The endpoints automatically communicate changes in service state to the call agent.

MGCP transactions are composed of a command and a mandatory response. There are eight types of commands:

- CreateConnection
- ModifyConnection
- DeleteConnection
- NotificationRequest
- Notify
- AuditEndpoint
- AuditConnection
- RestartInProgress

The first four commands are sent by the call agent to the gateway. The Notify command is sent by the gateway to the call agent. The gateway may also send a DeleteConnection. The registration of the MGCP gateway with the call agent is achieved by the RestartInProgress command. The AuditEndpoint and the AuditConnection commands are sent by the call agent to the gateway.

All commands are composed of a Command header, optionally followed by a session description. All responses are composed of a Response header, optionally followed by a session description.

- The port on which the gateway receives commands from the call agent. Gateways usually listen to UDP port 2427.

- The port on which the call agent receives commands from the gateway. Call agents usually listen to UDP port 2727.

**Note**

MGCP inspection does not support the use of different IP addresses for MGCP signaling and RTP data. A common and recommended practice is to send RTP data from a resilient IP address, such as a loopback or virtual IP address; however, the security appliance requires the RTP data to come from the same address as MGCP signalling.

## Configuring an MGCP Inspection Policy Map for Additional Inspection Control

If the network has multiple call agents and gateways for which the security appliance has to open pinholes, create an MGCP map. You can then apply the MGCP map when you enable MGCP inspection according to the [“Configuring Application Inspection” section on page 25-5](#)

To create an MGCP map, perform the following steps:

- Step 1** To create an MGCP inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect mgcp map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

- Step 2** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

- Step 3** To configure parameters that affect the inspection engine, perform the following steps:

- a.** To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap)# parameters
hostname(config-pmap-p)#
```

- b.** To configure the call agents, enter the following command for each call agent:

```
hostname(config-pmap-p)# call-agent ip_address group_id
```

Use the **call-agent** command to specify a group of call agents that can manage one or more gateways. The call agent group information is used to open connections for the call agents in the group (other than the one a gateway sends a command to) so that any of the call agents can send the response. call agents with the same *group\_id* belong to the same group. A call agent may belong to more than one group. The *group\_id* option is a number from 0 to 4294967295. The *ip\_address* option specifies the IP address of the call agent.

**Note**

MGCP call agents send AUEP messages to determine if MGCP end points are present. This establishes a flow through the security appliance and allows MGCP end points to register with the call agent.

- c.** To configure the gateways, enter the following command for each gateway:

```
hostname(config-pmap-p)# gateway ip_address group_id
```

Use the **gateway** command to specify which group of call agents are managing a particular gateway. The IP address of the gateway is specified with the *ip\_address* option. The *group\_id* option is a number from 0 to 4294967295 that must correspond with the *group\_id* of the call agents that are managing the gateway. A gateway may only belong to one group.

- d. If you want to change the maximum number of commands allowed in the MGCP command queue, enter the following command:

```
hostname(config-pmap-p)# command-queue command_limit
```

The following example shows how to define an MGCP map:

```
hostname(config)# policy-map type inspect mgcp sample_map
hostname(config-pmap)# parameters
hostname(config-pmap-p)# call-agent 10.10.11.5 101
hostname(config-pmap-p)# call-agent 10.10.11.6 101
hostname(config-pmap-p)# call-agent 10.10.11.7 102
hostname(config-pmap-p)# call-agent 10.10.11.8 102
hostname(config-pmap-p)# gateway 10.10.10.115 101
hostname(config-pmap-p)# gateway 10.10.10.116 102
hostname(config-pmap-p)# gateway 10.10.10.117 102
hostname(config-pmap-p)# command-queue 150
```

## Configuring MGCP Timeout Values

The **timeout mgcp command** lets you set the interval for inactivity after which an MGCP media connection is closed. The default is 5 minutes.

The **timeout mgcp-pat** command lets you set the timeout for PAT xlates. Because MGCP does not have a keepalive mechanism, if you use non-Cisco MGCP gateways (call agents), the PAT xlates are torn down after the default timeout interval, which is 30 seconds.

## Verifying and Monitoring MGCP Inspection

The **show mgcp commands** command lists the number of MGCP commands in the command queue. The **show mgcp sessions** command lists the number of existing MGCP sessions. The **detail** option includes additional information about each command (or session) in the output. The following is sample output from the **show mgcp commands** command:

```
hostname# show mgcp commands
1 in use, 1 most used, 200 maximum allowed
CRCX, gateway IP: host-pc-2, transaction ID: 2052, idle: 0:00:07
```

The following is sample output from the **show mgcp detail** command.

```
hostname# show mgcp commands detail
1 in use, 1 most used, 200 maximum allowed
CRCX, idle: 0:00:10
 Gateway IP host-pc-2
 Transaction ID 2052
 Endpoint name aaln/1
 Call ID 9876543210abcdef
 Connection ID
 Media IP 192.168.5.7
 Media port 6058
```

The following is sample output from the **show mgcp sessions** command.

```
hostname# show mgcp sessions
1 in use, 1 most used
Gateway IP host-pc-2, connection ID 6789af54c9, active 0:00:11
```

The following is sample output from the **show mgcp sessions detail** command.

```
hostname# show mgcp sessions detail
1 in use, 1 most used
Session active 0:00:14
 Gateway IP host-pc-2
 Call ID 9876543210abcdef
 Connection ID 6789af54c9
 Endpoint name aaln/1
 Media lcl port 6166
 Media rmt IP 192.168.5.7
 Media rmt port 6058
```

## NetBIOS Inspection

NetBIOS inspection is enabled by default. The NetBios inspection engine translates IP addresses in the NetBios name service (NBNS) packets according to the security appliance NAT configuration.

## Configuring a NetBIOS Inspection Policy Map for Additional Inspection Control

To specify actions when a message violates a parameter, create a NETBIOS inspection policy map. You can then apply the inspection policy map when you enable NETBIOS inspection according to the [“Configuring Application Inspection” section on page 25-5](#).

To create a NETBIOS inspection policy map, perform the following steps:

---

**Step 1** (Optional) Add one or more regular expressions for use in traffic matching commands according to the [“Creating a Regular Expression” section on page 21-12](#). See the types of text you can match in the **match** commands described in [Step 3](#).

**Step 2** (Optional) Create one or more regular expression class maps to group regular expressions according to the [“Creating a Regular Expression Class Map” section on page 21-14](#).

**Step 3** Create a NetBIOS inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect netbios policy_map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

**Step 4** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

**Step 5** To apply actions to matching traffic, perform the following steps.

- a. Specify the traffic on which you want to perform actions using one of the following methods:
  - Specify the NetBIOS class map that you created in [Step 3](#) by entering the following command:

```
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

- Specify traffic directly in the policy map using one of the **match** commands described in [Step 3](#). If you use a **match not** command, then any traffic that does not match the criterion in the **match not** command has the action applied.
- b. Specify the action you want to perform on the matching traffic by entering the following command:

```
hostname(config-pmap-c)# {[drop [send-protocol-error] |
drop-connection [send-protocol-error] | mask | reset] [log] | rate-limit message_rate}
```

Not all options are available for each **match** or **class** command. See the CLI help or the *Cisco Security Appliance Command Reference* for the exact options available.

The **drop** keyword drops all packets that match.

The **send-protocol-error** keyword sends a protocol error message.

The **drop-connection** keyword drops the packet and closes the connection.

The **mask** keyword masks out the matching portion of the packet.

The **reset** keyword drops the packet, closes the connection, and sends a TCP reset to the server and/or client.

The **log** keyword, which you can use alone or with one of the other keywords, sends a system log message.

The **rate-limit message\_rate** argument limits the rate of messages.

You can specify multiple **class** or **match** commands in the policy map. For information about the order of **class** and **match** commands, see the “[Defining Actions in an Inspection Policy Map](#)” section on [page 21-8](#).

**Step 6** To configure parameters that affect the inspection engine, perform the following steps:

- To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap)# parameters
hostname(config-pmap-p)#
```

- To check for NETBIOS protocol violations, enter the following command:

```
hostname(config-pmap-p)# protocol-violation [action [drop-connection / reset / log]]
```

Where the **drop-connection** action closes the connection. The **reset** action closes the connection and sends a TCP reset to the client. The **log** action sends a system log message when this policy map matches traffic.

The following example shows how to define a NETBIOS inspection policy map.

```
hostname(config)# policy-map type inspect netbios netbios_map
hostname(config-pmap)# protocol-violation drop log

hostname(config)# policy-map netbios_policy
hostname(config-pmap)# class inspection_default
hostname(config-pmap-c)# inspect netbios netbios_map
```

## PPTP Inspection

PPTP is a protocol for tunneling PPP traffic. A PPTP session is composed of one TCP channel and usually two PPTP GRE tunnels. The TCP channel is the control channel used for negotiating and managing the PPTP GRE tunnels. The GRE tunnels carries PPP sessions between the two hosts.

When enabled, PPTP application inspection inspects PPTP protocol packets and dynamically creates the GRE connections and xlates necessary to permit PPTP traffic. Only Version 1, as defined in RFC 2637, is supported.

PAT is only performed for the modified version of GRE [RFC 2637] when negotiated over the PPTP TCP control channel. Port Address Translation is *not* performed for the unmodified version of GRE [RFC 1701, RFC 1702].

Specifically, the security appliance inspects the PPTP version announcements and the outgoing call request/response sequence. Only PPTP Version 1, as defined in RFC 2637, is inspected. Further inspection on the TCP control channel is disabled if the version announced by either side is not Version 1. In addition, the outgoing-call request and reply sequence are tracked. Connections and xlates are dynamic allocated as necessary to permit subsequent secondary GRE data traffic.

The PPTP inspection engine must be enabled for PPTP traffic to be translated by PAT. Additionally, PAT is only performed for a modified version of GRE (RFC2637) and only if it is negotiated over the PPTP TCP control channel. PAT is not performed for the unmodified version of GRE (RFC 1701 and RFC 1702).

As described in RFC 2637, the PPTP protocol is mainly used for the tunneling of PPP sessions initiated from a modem bank PAC (PPTP Access Concentrator) to the headend PNS (PPTP Network Server). When used this way, the PAC is the remote client and the PNS is the server.

However, when used for VPN by Windows, the interaction is inverted. The PNS is a remote single-user PC that initiates connection to the head-end PAC to gain access to a central network.

## RADIUS Accounting Inspection

One of the well known problems is the over-billing attack in GPRS networks. The over-billing attack can cause consumers anger and frustration by being billed for services that they have not used. In this case, a malicious attacker sets up a connection to a server and obtains an IP address from the SGSN. When the attacker ends the call, the malicious server will still send packets to it, which gets dropped by the GGSN, but the connection from the server remains active. The IP address assigned to the malicious attacker gets released and reassigned to a legitimate user who will then get billed for services that the attacker will use.

RADIUS accounting inspection prevents this type of attack using by ensuring the traffic seen by the GGSN is legitimate. With the RADIUS accounting feature properly configured, the security appliance tears down a connection based on matching the Framed IP attribute in the Radius Accounting Request Start message with the Radius Accounting Request Stop message. When the Stop message is seen with the matching IP address in the Framed IP attribute, the security appliance looks for all connections with the source matching the IP address.

You have the option to configure a secret pre-shared key with the RADIUS server so the security appliance can validate the message. If the shared secret is not configured, the security appliance does not need to validate the source of the message and will only check that the source IP address is one of the configured addresses allowed to send the RADIUS messages.



## Configuring a RADIUS Inspection Policy Map for Additional Inspection Control

In order to use this feature, the **radius-accounting-map** will need to be specified in the **policy-map** and then applied to the service-policy to specify that this traffic is for to-the-box inspection.

The following example shows the complete set of commands in context to properly configure this feature:

---

**Step 1** Configure the class map and the port:

```
class-map type management c1
 match port udp eq 1813
```

**Step 2** Create the policy map, and configure the parameters for RADIUS accounting inspection using the parameter command to access the proper mode to configure the attributes, host, and key.

```
policy-map type inspect radius-accounting radius_accounting_map
 parameters
 host 10.1.1.1 inside key 123456789
 send response
 enable gprs
 validate-attribute 31
```

**Step 3** Configure the service policy.

```
policy-map global_policy
 class c1
 inspect radius-accounting radius_accounting_map

service-policy global_policy global
```

## RSH Inspection

RSH inspection is enabled by default. The RSH protocol uses a TCP connection from the RSH client to the RSH server on TCP port 514. The client and server negotiate the TCP port number where the client listens for the STDERR output stream. RSH inspection supports NAT of the negotiated port number if necessary.

## RTSP Inspection

This section describes RTSP application inspection. This section includes the following topics:

- [RTSP Inspection Overview, page 25-63](#)
- [Using RealPlayer, page 25-64](#)
- [Restrictions and Limitations, page 25-64](#)

## RTSP Inspection Overview

The RTSP inspection engine lets the security appliance pass RTSP packets. RTSP is used by RealAudio, RealNetworks, Apple QuickTime 4, RealPlayer, and Cisco IP/TV connections.

**Note**

For Cisco IP/TV, use RTSP TCP port 554 and TCP 8554.

RTSP applications use the well-known port 554 with TCP (rarely UDP) as a control channel. The security appliance only supports TCP, in conformity with RFC 2326. This TCP control channel is used to negotiate the data channels that is used to transmit audio/video traffic, depending on the transport mode that is configured on the client.

The supported RDT transports are: rtp/avp, rtp/avp/udp, x-real-rdt, x-real-rdt/udp, and x-pn-tng/udp.

The security appliance parses Setup response messages with a status code of 200. If the response message is travelling inbound, the server is outside relative to the security appliance and dynamic channels need to be opened for connections coming inbound from the server. If the response message is outbound, then the security appliance does not need to open dynamic channels.

Because RFC 2326 does not require that the client and server ports must be in the SETUP response message, the security appliance keeps state and remembers the client ports in the SETUP message. QuickTime places the client ports in the SETUP message and then the server responds with only the server ports.

RTSP inspection does not support PAT or dual-NAT. Also, the security appliance cannot recognize HTTP cloaking where RTSP messages are hidden in the HTTP messages.

## Using RealPlayer

When using RealPlayer, it is important to properly configure transport mode. For the security appliance, add an **access-list** command from the server to the client or vice versa. For RealPlayer, change transport mode by clicking **Options>Preferences>Transport>RTSP Settings**.

If using TCP mode on the RealPlayer, select the **Use TCP to Connect to Server** and **Attempt to use TCP for all content** check boxes. On the security appliance, there is no need to configure the inspection engine.

If using UDP mode on the RealPlayer, select the **Use TCP to Connect to Server** and **Attempt to use UDP for static content** check boxes, and for live content not available via Multicast. On the security appliance, add an **inspect rtsp port** command.

## Restrictions and Limitations

The following restrictions apply to the **inspect rtsp** command.

- The security appliance does not support multicast RTSP or RTSP messages over UDP.
- PAT is not supported.
- The security appliance does not have the ability to recognize HTTP cloaking where RTSP messages are hidden in the HTTP messages.
- The security appliance cannot perform NAT on RTSP messages because the embedded IP addresses are contained in the SDP files as part of HTTP or RTSP messages. Packets could be fragmented and security appliance cannot perform NAT on fragmented packets.
- With Cisco IP/TV, the number of translates the security appliance performs on the SDP part of the message is proportional to the number of program listings in the Content Manager (each program listing can have at least six embedded IP addresses).
- You can configure NAT for Apple QuickTime 4 or RealPlayer. Cisco IP/TV only works with NAT if the Viewer and Content Manager are on the outside network and the server is on the inside network.

# SIP Inspection

This section describes SIP application inspection. This section includes the following topics:

- [SIP Inspection Overview, page 25-65](#)
- [SIP Instant Messaging, page 25-65](#)
- [Configuring SIP Timeout Values, page 25-70](#)
- [Verifying and Monitoring SIP Inspection, page 25-70](#)

## SIP Inspection Overview

SIP, as defined by the IETF, enables call handling sessions, particularly two-party audio conferences, or “calls.” SIP works with SDP for call signalling. SDP specifies the ports for the media stream. Using SIP, the security appliance can support any SIP VoIP gateways and VoIP proxy servers. SIP and SDP are defined in the following RFCs:

- SIP: Session Initiation Protocol, RFC 3261
- SDP: Session Description Protocol, RFC 2327

To support SIP calls through the security appliance, signaling messages for the media connection addresses, media ports, and embryonic connections for the media must be inspected, because while the signaling is sent over a well-known destination port (UDP/TCP 5060), the media streams are dynamically allocated. Also, SIP embeds IP addresses in the user-data portion of the IP packet. SIP inspection applies NAT for these embedded IP addresses.

The following limitations and restrictions apply when using PAT with SIP:

- If a remote endpoint tries to register with a SIP proxy on a network protected by the security appliance, the registration fails under very specific conditions, as follows:
  - PAT is configured for the remote endpoint.
  - The SIP registrar server is on the outside network.
  - The port is missing in the contact field in the REGISTER message sent by the endpoint to the proxy server.
- If a SIP device transmits a packet in which the SDP portion has an IP address in the owner/creator field (o=) that is different than the IP address in the connection field (c=), the IP address in the o= field may not be properly translated. This is due to a limitation in the SIP protocol, which does not provide a port value in the o= field.

## SIP Instant Messaging

Instant Messaging refers to the transfer of messages between users in near real-time. SIP supports the Chat feature on Windows XP using Windows Messenger RTC Client version 4.7.0105 only. The MESSAGE/INFO methods and 202 Accept response are used to support IM as defined in the following RFCs:

- Session Initiation Protocol (SIP)-Specific Event Notification, RFC 3265
- Session Initiation Protocol (SIP) Extension for Instant Messaging, RFC 3428

MESSAGE/INFO requests can come in at any time after registration/subscription. For example, two users can be online at any time, but not chat for hours. Therefore, the SIP inspection engine opens pinholes that time out according to the configured SIP timeout value. This value must be configured at least five minutes longer than the subscription duration. The subscription duration is defined in the Contact Expires value and is typically 30 minutes.

Because MESSAGE/INFO requests are typically sent using a dynamically allocated port other than port 5060, they are required to go through the SIP inspection engine.

**Note**

Only the Chat feature is currently supported. Whiteboard, File Transfer, and Application Sharing are not supported. RTC Client 5.0 is not supported.

SIP inspection translates the SIP text-based messages, recalculates the content length for the SDP portion of the message, and recalculates the packet length and checksum. It dynamically opens media connections for ports specified in the SDP portion of the SIP message as address/ports on which the endpoint should listen.

SIP inspection has a database with indices CALL\_ID/FROM/TO from the SIP payload. These indices identify the call, the source, and the destination. This database contains the media addresses and media ports found in the SDP media information fields and the media type. There can be multiple media addresses and ports for a session. The security appliance opens RTP/RTCP connections between the two endpoints using these media addresses/ports.

The well-known port 5060 must be used on the initial call setup (INVITE) message; however, subsequent messages may not have this port number. The SIP inspection engine opens signaling connection pinholes, and marks these connections as SIP connections. This is done for the messages to reach the SIP application and be translated.

As a call is set up, the SIP session is in the “transient” state until the media address and media port is received from the called endpoint in a Response message indicating the RTP port the called endpoint listens on. If there is a failure to receive the response messages within one minute, the signaling connection is torn down.

Once the final handshake is made, the call state is moved to active and the signaling connection remains until a BYE message is received.

If an inside endpoint initiates a call to an outside endpoint, a media hole is opened to the outside interface to allow RTP/RTCP UDP packets to flow to the inside endpoint media address and media port specified in the INVITE message from the inside endpoint. Unsolicited RTP/RTCP UDP packets to an inside interface does not traverse the security appliance, unless the security appliance configuration specifically allows it.

## Configuring a SIP Inspection Policy Map for Additional Inspection Control

To specify actions when a message violates a parameter, create a SIP inspection policy map. You can then apply the inspection policy map when you enable SIP inspection according to the [“Configuring Application Inspection” section on page 25-5](#).

To create a SIP inspection policy map, perform the following steps:

- Step 1** (Optional) Add one or more regular expressions for use in traffic matching commands according to the [“Creating a Regular Expression” section on page 21-12](#). See the types of text you can match in the **match** commands described in [Step 3](#).

**Step 2** (Optional) Create one or more regular expression class maps to group regular expressions according to the “[Creating a Regular Expression Class Map](#)” section on page 21-14.s

**Step 3** (Optional) Create a SIP inspection class map by performing the following steps.

A class map groups multiple traffic matches. Traffic must match *all* of the **match** commands to match the class map. You can alternatively identify **match** commands directly in the policy map. The difference between creating a class map and defining the traffic match directly in the inspection policy map is that the class map lets you create more complex match criteria, and you can reuse class maps.

To specify traffic that should not match the class map, use the **match not** command. For example, if the **match not** command specifies the string “example.com,” then any traffic that includes “example.com” does not match the class map.

For the traffic that you identify in this class map, you can specify actions such as drop-connection, reset, and/or log the connection in the inspection policy map.

If you want to perform different actions for each **match** command, you should identify the traffic directly in the policy map.

- a. Create the class map by entering the following command:

```
hostname(config)# class-map type inspect sip [match-all] class_map_name
hostname(config-cmap)#
```

Where *the class\_map\_name* is the name of the class map. The **match-all** keyword specifies that traffic must match all criteria to match the class map. **match-all** is the default and only option. The CLI enters class-map configuration mode, where you can enter one or more **match** commands.

- b. (Optional) To add a description to the class map, enter the following command:

```
hostname(config-cmap)# description string
```

Where *string* is the description of the class map (up to 200 characters).

- c. (Optional) To match a called party, as specified in the To header, enter the following command:

```
hostname(config-cmap)# match [not] called-party regex {class class_name | regex_name}
```

Where the **regex regex\_name** argument is the regular expression you created in [Step 1](#). The **class regex\_class\_name** is the regular expression class map you created in [Step 2](#).

- d. (Optional) To match a calling party, as specified in the From header, enter the following command:

```
hostname(config-cmap)# match [not] calling-party regex {class class_name | regex_name}
```

Where the **regex regex\_name** argument is the regular expression you created in [Step 1](#). The **class regex\_class\_name** is the regular expression class map you created in [Step 2](#).

- e. (Optional) To match a content length in the SIP header, enter the following command:

```
hostname(config-cmap)# match [not] content length gt length
```

Where *length* is the number of bytes the content length is greater than. 0 to 65536.

- f. (Optional) To match an SDP content type or regular expression, enter the following command:

```
hostname(config-cmap)# match [not] content type {sdp | regex {class class_name | regex_name}}
```

Where the **regex regex\_name** argument is the regular expression you created in [Step 1](#). The **class regex\_class\_name** is the regular expression class map you created in [Step 2](#).

- g. (Optional) To match a SIP IM subscriber, enter the following command:

```
hostname(config-cmap)# match [not] im-subscriber regex {class class_name | regex_name}
```

Where the **regex** *regex\_name* argument is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#).

- h. (Optional) To match a SIP via header, enter the following command:

```
hostname(config-cmap)# match [not] message-path regex {class class_name | regex_name}
```

Where the **regex** *regex\_name* argument is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#).

- i. (Optional) To match a SIP request method, enter the following command:

```
hostname(config-cmap)# match [not] request-method method
```

Where *method* is the type of method to match (ack, bye, cancel, info, invite, message, notify, options, prack, refer, register, subscribe, unknown, update).

- j. (Optional) To match the requester of a third-party registration, enter the following command:

```
hostname(config-cmap)# match [not] third-party-registration regex {class class_name | regex_name}
```

Where the **regex** *regex\_name* argument is the regular expression you created in [Step 1](#). The **class** *regex\_class\_name* is the regular expression class map you created in [Step 2](#).

- k. (Optional) To match an URI in the SIP headers, enter the following command:

```
hostname(config-cmap)# match [not] uri {sip | tel} length gt length
```

Where *length* is the number of bytes the URI is greater than. 0 to 65536.

- Step 4** Create a SIP inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect sip policy_map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

- Step 5** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

- Step 6** To apply actions to matching traffic, perform the following steps.

- a. Specify the traffic on which you want to perform actions using one of the following methods:

- Specify the SIP class map that you created in [Step 3](#) by entering the following command:

```
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

- Specify traffic directly in the policy map using one of the **match** commands described in [Step 3](#). If you use a **match not** command, then any traffic that does not match the criterion in the **match not** command has the action applied.

- b. Specify the action you want to perform on the matching traffic by entering the following command:

```
hostname(config-pmap-c)# {[drop [send-protocol-error] |
drop-connection [send-protocol-error] | mask | reset] [log] | rate-limit message_rate}
```

Not all options are available for each **match** or **class** command. See the CLI help or the *Cisco Security Appliance Command Reference* for the exact options available.

The **drop** keyword drops all packets that match.

The **send-protocol-error** keyword sends a protocol error message.

The **drop-connection** keyword drops the packet and closes the connection.

The **mask** keyword masks out the matching portion of the packet.

The **reset** keyword drops the packet, closes the connection, and sends a TCP reset to the server and/or client.

The **log** keyword, which you can use alone or with one of the other keywords, sends a system log message.

The **rate-limit** *message\_rate* argument limits the rate of messages.

You can specify multiple **class** or **match** commands in the policy map. For information about the order of **class** and **match** commands, see the “[Defining Actions in an Inspection Policy Map](#)” section on page 21-8.

**Step 7** To configure parameters that affect the inspection engine, perform the following steps:

- a. To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap) # parameters
hostname(config-pmap-p) #
```

- b. To enable or disable instant messaging, enter the following command:

```
hostname(config-pmap-p) # im
```

- c. To enable or disable IP address privacy, enter the following command:

```
hostname(config-pmap-p) # ip-address-privacy
```

- d. To enable check on Max-forwards header field being 0 (which cannot be 0 before reaching the destination), enter the following command:

```
hostname(config-pmap-p) # max-forwards-validation action {drop | drop-connection |
reset | log} [log]
```

- e. To enable check on RTP packets flowing on the pinholes for protocol conformance, enter the following command:

```
hostname(config-pmap-p) # rtp-conformance [enforce-payloadtype]
```

Where the **enforce-payloadtype** keyword enforces the payload type to be audio or video based on the signaling exchange.

- f. To identify the Server and User-Agent header fields, which expose the software version of either a server or an endpoint, enter the following command:

```
hostname(config-pmap-p) # software-version action {mask | log} [log]
```

Where the **mask** keyword masks the software version in the SIP messages.

- g. To enable state checking validation, enter the following command:

```
hostname(config-pmap-p) # state-checking action {drop | drop-connection | reset | log}
[log]
```

- h. To enable strict verification of the header fields in the SIP messages according to RFC 3261, enter the following command:

```
hostname(config-pmap-p) # strict-header-validation action {drop | drop-connection |
reset | log} [log]
```

- i. To allow non SIP traffic using the well-known SIP signaling port, enter the following command:

```
hostname(config-pmap-p) # traffic-non-sip
```

- j. To identify the non-SIP URIs present in the Alert-Info and Call-Info header fields, enter the following command:

```
hostname(config-pmap-p)# uri-non-sip action {mask | log} [log]
```

The following example shows how to disable instant messaging over SIP:

```
hostname(config)# policy-map type inspect sip mymap
hostname(config-pmap)# parameters
hostname(config-pmap-p)# no im

hostname(config)# policy-map global_policy
hostname(config-pmap)# class inspection_default
hostname(config-pmap-c)# inspect sip mymap

hostname(config)# service-policy global_policy global
```

## Configuring SIP Timeout Values

The media connections are torn down within two minutes after the connection becomes idle. This is, however, a configurable timeout and can be set for a shorter or longer period of time. To configure the timeout for the SIP control connection, enter the following command:

```
hostname(config)# timeout sip hh:mm:ss
```

This command configures the idle timeout after which a SIP control connection is closed.

To configure the timeout for the SIP media connection, enter the following command:

```
hostname(config)# timeout sip_media hh:mm:ss
```

This command configures the idle timeout after which a SIP media connection is closed.

## Verifying and Monitoring SIP Inspection

The **show sip** command assists in troubleshooting SIP inspection engine issues and is described with the **inspect protocol sip udp 5060** command. The **show timeout sip** command displays the timeout value of the designated protocol.

The **show sip** command displays information for SIP sessions established across the security appliance. Along with the **debug sip** and **show local-host** commands, this command is used for troubleshooting SIP inspection engine issues.



### Note

We recommend that you configure the **pager** command before entering the **show sip** command. If there are a lot of SIP session records and the **pager** command is not configured, it takes a while for the **show sip** command output to reach its end.

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

```
hostname# show sip
Total: 2
call-id c3943000-960ca-2e43-228f@10.130.56.44
state Call init, idle 0:00:01
call-id c3943000-860ca-7e1f-11f7@10.130.56.45
state Active, idle 0:00:06
```



This sample shows two active SIP sessions on the security appliance (as shown in the Total field). Each call-id represents a call.

The first session, with the call-id c3943000-960ca-2e43-228f@10.130.56.44, is in the state Call Init, which means the session is still in call setup. Call setup is not complete until a final response to the call has been received. For instance, the caller has already sent the INVITE, and maybe received a 100 Response, but has not yet seen the 200 OK, so the call setup is not complete yet. Any non-1xx response message is considered a final response. This session has been idle for 1 second.

The second session is in the state Active, in which call setup is complete and the endpoints are exchanging media. This session has been idle for 6 seconds.

## Skinny (SCCP) Inspection

This section describes SCCP application inspection. This section includes the following topics:

- [SCCP Inspection Overview, page 25-71](#)
- [Supporting Cisco IP Phones, page 25-71](#)
- [Restrictions and Limitations, page 25-72](#)
- [Verifying and Monitoring SCCP Inspection, page 25-72](#)

## SCCP Inspection Overview

Skinny (SCCP) is a simplified protocol used in VoIP networks. Cisco IP Phones using SCCP can coexist in an H.323 environment. When used with Cisco CallManager, the SCCP client can interoperate with H.323 compliant terminals. Application layer functions in the security appliance recognize SCCP Version 3.3. There are 5 versions of the SCCP protocol: 2.4, 3.0.4, 3.1.1, 3.2, and 3.3.2. The security appliance supports all versions through Version 3.3.2.

The security appliance supports PAT and NAT for SCCP. PAT is necessary if you have more IP phones than global IP addresses for the IP phones to use. By supporting NAT and PAT of SCCP Signaling packets, Skinny application inspection ensures that all SCCP signalling and media packets can traverse the security appliance.

Normal traffic between Cisco CallManager and Cisco IP Phones uses SCCP and is handled by SCCP inspection without any special configuration. The security appliance also supports DHCP options 150 and 66, which it accomplishes by sending the location of a TFTP server to Cisco IP Phones and other DHCP clients. Cisco IP Phones might also include DHCP option 3 in their requests, which sets the default route. For more information, see the [“Using Cisco IP Phones with a DHCP Server” section on page 10-4](#).

## Supporting Cisco IP Phones

In topologies where Cisco CallManager is located on the higher security interface with respect to the Cisco IP Phones, if NAT is required for the Cisco CallManager IP address, the mapping must be **static** as a Cisco IP Phone requires the Cisco CallManager IP address to be specified explicitly in its configuration. An static identity entry allows the Cisco CallManager on the higher security interface to accept registrations from the Cisco IP Phones.

Cisco IP Phones require access to a TFTP server to download the configuration information they need to connect to the Cisco CallManager server.

When the Cisco IP Phones are on a lower security interface compared to the TFTP server, you must use an access list to connect to the protected TFTP server on UDP port 69. While you do need a static entry for the TFTP server, this does not have to be an identity static entry. When using NAT, an identity static entry maps to the same IP address. When using PAT, it maps to the same IP address and port.

When the Cisco IP Phones are on a *higher* security interface compared to the TFTP server and Cisco CallManager, no access list or static entry is required to allow the Cisco IP Phones to initiate the connection.

## Restrictions and Limitations

The following are limitations that apply to the current version of PAT and NAT support for SCCP:

- PAT does not work with configurations containing the **alias** command.
- Outside NAT or PAT is *not* supported.

If the address of an internal Cisco CallManager is configured for NAT or PAT to a different IP address or port, registrations for external Cisco IP Phones fail because the security appliance currently does not support NAT or PAT for the file content transferred over TFTP. Although the security appliance supports NAT of TFTP messages and opens a pinhole for the TFTP file, the security appliance cannot translate the Cisco CallManager IP address and port embedded in the Cisco IP Phone configuration files that are transferred by TFTP during phone registration.



### Note

The security appliance supports stateful failover of SCCP calls except for calls that are in the middle of call setup.

## Verifying and Monitoring SCCP Inspection

The **show skinny** command assists in troubleshooting SCCP (Skinny) inspection engine issues. The following is sample output from the **show skinny** command under the following conditions. There are two active Skinny sessions set up across the security appliance. The first one is established between an internal Cisco IP Phone at local address 10.0.0.11 and an external Cisco CallManager at 172.18.1.33. TCP port 2000 is the CallManager. The second one is established between another internal Cisco IP Phone at local address 10.0.0.22 and the same Cisco CallManager.

```
hostname# show skinny
 LOCAL FOREIGN STATE

1 10.0.0.11/52238 172.18.1.33/2000 1
 MEDIA 10.0.0.11/22948 172.18.1.22/20798
2 10.0.0.22/52232 172.18.1.33/2000 1
 MEDIA 10.0.0.22/20798 172.18.1.11/22948
```

The output indicates that a call has been established between two internal Cisco IP Phones. The RTP listening ports of the first and second phones are UDP 22948 and 20798 respectively.

The following is sample output from the **show xlate debug** command for these Skinny connections:

```
hostname# show xlate debug
2 in use, 2 most used
Flags: D - DNS, d - dump, I - identity, i - inside, n - no random,
 r - portmap, s - static
NAT from inside:10.0.0.11 to outside:172.18.1.11 flags si idle 0:00:16 timeout 0:05:00
NAT from inside:10.0.0.22 to outside:172.18.1.22 flags si idle 0:00:14 timeout 0:05:00
```

## Configuring a Skinny (SCCP) Inspection Policy Map for Additional Inspection Control

To specify actions when a message violates a parameter, create an SCCP inspection policy map. You can then apply the inspection policy map when you enable SCCP inspection according to the [“Configuring Application Inspection” section on page 25-5](#).

To create an SCCP inspection policy map, perform the following steps:

**Step 1** (Optional) Add one or more regular expressions for use in traffic matching commands according to the [“Creating a Regular Expression” section on page 21-12](#). See the types of text you can match in the **match** commands described in [Step 3](#).

**Step 2** (Optional) Create one or more regular expression class maps to group regular expressions according to the [“Creating a Regular Expression Class Map” section on page 21-14](#).

**Step 3** Create an SCCP inspection policy map, enter the following command:

```
hostname(config)# policy-map type inspect skinny policy_map_name
hostname(config-pmap)#
```

Where the *policy\_map\_name* is the name of the policy map. The CLI enters policy-map configuration mode.

**Step 4** (Optional) To add a description to the policy map, enter the following command:

```
hostname(config-pmap)# description string
```

**Step 5** To apply actions to matching traffic, perform the following steps.

a. Specify the traffic on which you want to perform actions using one of the following methods:

- Specify the SCCP class map that you created in [Step 3](#) by entering the following command:

```
hostname(config-pmap)# class class_map_name
hostname(config-pmap-c)#
```

- Specify traffic directly in the policy map using one of the **match** commands described in [Step 3](#). If you use a **match not** command, then any traffic that does not match the criterion in the **match not** command has the action applied.

b. Specify the action you want to perform on the matching traffic by entering the following command:

```
hostname(config-pmap-c)# {[drop [send-protocol-error] |
drop-connection [send-protocol-error] | mask | reset] [log] | rate-limit message_rate}
```

Not all options are available for each **match** or **class** command. See the CLI help or the *Cisco Security Appliance Command Reference* for the exact options available.

The **drop** keyword drops all packets that match.

The **send-protocol-error** keyword sends a protocol error message.

The **drop-connection** keyword drops the packet and closes the connection.

The **mask** keyword masks out the matching portion of the packet.

The **reset** keyword drops the packet, closes the connection, and sends a TCP reset to the server and/or client.

The **log** keyword, which you can use alone or with one of the other keywords, sends a system log message.

The **rate-limit** *message\_rate* argument limits the rate of messages.

**Step 6** You can specify multiple **class** or **match** commands in the policy map. For information about the order of **class** and **match** commands, see the [“Defining Actions in an Inspection Policy Map” section on page 21-8](#). To configure parameters that affect the inspection engine, perform the following steps:

- a. To enter parameters configuration mode, enter the following command:

```
hostname(config-pmap)# parameters
hostname(config-pmap-p)#
```

- b. To enforce registration before calls can be placed, enter the following command:

```
hostname(config-pmap-p)# enforce-registration
```

- c. To set the maximum SCCP station message ID allowed, enter the following command:

```
hostname(config-pmap-p)# message-ID max hex_value
```

Where the *hex\_value* argument is the station message ID in hex.

- d. To check RTP packets flowing on the pinholes for protocol conformance, enter the following command:

```
hostname(config-pmap-p)# rtp-conformance [enforce-payloadtype]
```

Where the **enforce-payloadtype** keyword enforces the payload type to be audio or video based on the signaling exchange.

- e. To set the maximum and minimum SCCP prefix length value allowed, enter the following command:

```
hostname(config-pmap-p)# sccp-prefix-len {max | min} value_length
```

Where the *value\_length* argument is a maximum or minimum value.

- f. To configure the timeout value for signaling and media connections, enter the following command:

```
hostname(config-pmap-p)# timeout
```

The following example shows how to define an SCCP inspection policy map.

```
hostname(config)# policy-map type inspect skinny skinny-map
hostname(config-pmap)# parameters
hostname(config-pmap-p)# enforce-registration
hostname(config-pmap-p)# match message-id range 200 300
hostname(config-pmap-p)# drop log
hostname(config)# class-map inspection_default
hostname(config-cmap)# match default-inspection-traffic
hostname(config)# policy-map global_policy
hostname(config-pmap)# class inspection_default
hostname(config-pmap-c)# inspect skinny skinny-map
hostname(config)# service-policy global_policy global
```

## SMTP and Extended SMTP Inspection

ESMTP application inspection provides improved protection against SMTP-based attacks by restricting the types of SMTP commands that can pass through the security appliance and by adding monitoring capabilities.

ESMTP is an enhancement to the SMTP protocol and is similar in most respects to SMTP. For convenience, the term SMTP is used in this document to refer to both SMTP and ESMTP. The application inspection process for extended SMTP is similar to SMTP application inspection and includes support for SMTP sessions. Most commands used in an extended SMTP session are the same as those used in an SMTP session but an ESMTP session is considerably faster and offers more options related to reliability and security, such as delivery status notification.

Extended SMTP application inspection adds support for eight extended SMTP commands, including AUTH, EHLO, ETRN, HELP, SAML, SEND, SOML and VRFY. Along with the support for seven RFC 821 commands (DATA, HELO, MAIL, NOOP, QUIT, RCPT, RSET), the security appliance supports a total of fifteen SMTP commands.

Other extended SMTP commands, such as ATRN, STARTLS, ONEX, VERB, CHUNKING, and private extensions are not supported. Unsupported commands are translated into Xs, which are rejected by the internal server. This results in a message such as “500 Command unknown: 'XXX'.” Incomplete commands are discarded.

The ESMTP inspection engine changes the characters in the server SMTP banner to asterisks except for the “2”, “0”, “0” characters. Carriage return (CR) and linefeed (LF) characters are ignored.

With SMTP inspection enabled, a Telnet session used for interactive SMTP may hang if the following rules are not observed: SMTP commands must be at least four characters in length; must be terminated with carriage return and line feed; and must wait for a response before issuing the next reply.

An SMTP server responds to client requests with numeric reply codes and optional human-readable strings. SMTP application inspection controls and reduces the commands that the user can use as well as the messages that the server returns. SMTP inspection performs three primary tasks:

- Restricts SMTP requests to seven basic SMTP commands and eight extended commands.
- Monitors the SMTP command-response sequence.
- Generates an audit trail—Audit record 108002 is generated when invalid character embedded in the mail address is replaced. For more information, see RFC 821.

SMTP inspection monitors the command and response sequence for the following anomalous signatures:

- Truncated commands.
- Incorrect command termination (not terminated with <CR><LR>).
- The MAIL and RCPT commands specify who are the sender and the receiver of the mail. Mail addresses are scanned for strange characters. The pipeline character (|) is deleted (changed to a blank space) and “<”, “>” are only allowed if they are used to define a mail address (“>” must be preceded by “<”).
- Unexpected transition by the SMTP server.
- For unknown commands, the security appliance changes all the characters in the packet to X. In this case, the server generates an error code to the client. Because of the change in the packet, the TCP checksum has to be recalculated or adjusted.
- TCP stream editing.
- Command pipelining.

# SNMP Inspection

SNMP application inspection lets you restrict SNMP traffic to a specific version of SNMP. Earlier versions of SNMP are less secure; therefore, denying certain SNMP versions may be required by your security policy. The security appliance can deny SNMP versions 1, 2, 2c, or 3. You control the versions permitted by creating an SNMP map. You then apply the SNMP map when you enable SNMP inspection according to the [“Configuring Application Inspection” section on page 25-5](#).

To create an SNMP inspection policy map, perform the following steps:

---

**Step 1** To create an SNMP map, enter the following command:

```
hostname(config)# snmp-map map_name
hostname(config-snmp-map)#
```

where *map\_name* is the name of the SNMP map. The CLI enters SNMP map configuration mode.

**Step 2** To specify the versions of SNMP to deny, enter the following command for each version:

```
hostname(config-snmp-map)# deny version version
hostname(config-snmp-map)#
```

where *version* is 1, 2, 2c, or 3.

---

The following example denies SNMP Versions 1 and 2:

```
hostname(config)# snmp-map sample_map
hostname(config-snmp-map)# deny version 1
hostname(config-snmp-map)# deny version 2
```

# SQL\*Net Inspection

SQL\*Net inspection is enabled by default.

The SQL\*Net protocol consists of different packet types that the security appliance handles to make the data stream appear consistent to the Oracle applications on either side of the security appliance.

The default port assignment for SQL\*Net is 1521. This is the value used by Oracle for SQL\*Net, but this value does not agree with IANA port assignments for Structured Query Language (SQL). Use the **class-map** command to apply SQL\*Net inspection to a range of port numbers.

The security appliance translates all addresses and looks in the packets for all embedded ports to open for SQL\*Net Version 1.

For SQL\*Net Version 2, all DATA or REDIRECT packets that immediately follow REDIRECT packets with a zero data length will be fixed up.

The packets that need fix-up contain embedded host/port addresses in the following format:

```
(ADDRESS=(PROTOCOL=tcp) (DEV=6) (HOST=a.b.c.d) (PORT=a))
```

SQL\*Net Version 2 TNSFrame types (Connect, Accept, Refuse, Resend, and Marker) will not be scanned for addresses to NAT nor will inspection open dynamic connections for any embedded ports in the packet.

SQL\*Net Version 2 TNSFrames, Redirect, and Data packets will be scanned for ports to open and addresses to NAT, if preceded by a REDIRECT TNSFrame type with a zero data length for the payload. When the Redirect message with data length zero passes through the security appliance, a flag will be set in the connection data Structure to expect the Data or Redirect message that follows to be translated and ports to be dynamically opened. If one of the TNS frames in the preceding paragraph arrive after the Redirect message, the flag will be reset.

The SQL\*Net inspection engine will recalculate the checksum, change IP, TCP lengths, and readjust Sequence Numbers and Acknowledgment Numbers using the delta of the length of the new and old message.

SQL\*Net Version 1 is assumed for all other cases. TNSFrame types (Connect, Accept, Refuse, Resend, Marker, Redirect, and Data) and all packets will be scanned for ports and addresses. Addresses will be translated and port connections will be opened.

## Sun RPC Inspection

This section describes Sun RPC application inspection. This section includes the following topics:

- [Sun RPC Inspection Overview, page 25-77](#)
- [Managing Sun RPC Services, page 25-77](#)
- [Verifying and Monitoring Sun RPC Inspection, page 25-78](#)

## Sun RPC Inspection Overview

The Sun RPC inspection engine enables or disables application inspection for the Sun RPC protocol. Sun RPC is used by NFS and NIS. Sun RPC services can run on any port. When a client attempts to access an Sun RPC service on a server, it must learn the port that service is running on. It does this by querying the port mapper process, usually rpcbind, on the well-known port of 111.

The client sends the Sun RPC program number of the service and the port mapper process responds with the port number of the service. The client sends its Sun RPC queries to the server, specifying the port identified by the port mapper process. When the server replies, the security appliance intercepts this packet and opens both embryonic TCP and UDP connections on that port.



### Note

NAT or PAT of Sun RPC payload information is not supported.

## Managing Sun RPC Services

Use the Sun RPC services table to control Sun RPC traffic through the security appliance based on established Sun RPC sessions. To create entries in the Sun RPC services table, use the **sunrpc-server** command in global configuration mode:

```
hostname(config)# sunrpc-server interface_name ip_address mask service service_type
protocol {tcp | udp} port[-port] timeout hh:mm:ss
```

You can use this command to specify the timeout after which the pinhole that was opened by Sun RPC application inspection will be closed. For example, to create a timeout of 30 minutes to the Sun RPC server with the IP address 192.168.100.2, enter the following command:

```
hostname(config)# sunrpc-server inside 192.168.100.2 255.255.255.255 service 100003
protocol tcp 111 timeout 00:30:00
```

This command specifies that the pinhole that was opened by Sun RPC application inspection will be closed after 30 minutes. In this example, the Sun RPC server is on the inside interface using TCP port 111. You can also specify UDP, a different port number, or a range of ports. To specify a range of ports, separate the starting and ending port numbers in the range with a hyphen (for example, 111-113).

The service type identifies the mapping between a specific service type and the port number used for the service. To determine the service type, which in this example is 100003, use the **sunrpcinfo** command at the UNIX or Linux command line on the Sun RPC server machine.

To clear the Sun RPC configuration, enter the following command.

```
hostname(config)# clear configure sunrpc-server
```

This removes the configuration performed using the **sunrpc-server** command. The **sunrpc-server** command allows pinholes to be created with a specified timeout.

To clear the active Sun RPC services, enter the following command:

```
hostname(config)# clear sunrpc-server active
```

This clears the pinholes that are opened by Sun RPC application inspection for specific services, such as NFS or NIS.

## Verifying and Monitoring Sun RPC Inspection

The sample output in this section is for a Sun RPC server with an IP address of 192.168.100.2 on the inside interface and a Sun RPC client with an IP address of 209.168.200.5 on the outside interface.

To view information about the current Sun RPC connections, enter the **show conn** command. The following is sample output from the **show conn** command:

```
hostname# show conn
15 in use, 21 most used
UDP out 209.165.200.5:800 in 192.168.100.2:2049 idle 0:00:04 flags -
UDP out 209.165.200.5:714 in 192.168.100.2:111 idle 0:00:04 flags -
UDP out 209.165.200.5:712 in 192.168.100.2:647 idle 0:00:05 flags -
UDP out 192.168.100.2:0 in 209.165.200.5:714 idle 0:00:05 flags i
hostname(config)#
```

To display the information about the Sun RPC service table configuration, enter the **show running-config sunrpc-server** command. The following is sample output from the **show running-config sunrpc-server** command:

```
hostname(config)# show running-config sunrpc-server
sunrpc-server inside 192.168.100.2 255.255.255.255 service 100003 protocol UDP port 111
timeout 0:30:00
sunrpc-server inside 192.168.100.2 255.255.255.255 service 100005 protocol UDP port 111
timeout 0:30:00
```

This output shows that a timeout interval of 30 minutes is configured on UDP port 111 for the Sun RPC server with the IP address 192.168.100.2 on the inside interface.

To display the pinholes open for Sun RPC services, enter the **show sunrpc-server active** command. The following is sample output from **show sunrpc-server active** command:

```
hostname# show sunrpc-server active
LOCAL FOREIGN SERVICE TIMEOUT

1 209.165.200.5/0 192.168.100.2/2049 100003 0:30:00
```



```

2 209.165.200.5/0 192.168.100.2/2049 100003 0:30:00
3 209.165.200.5/0 192.168.100.2/647 100005 0:30:00
4 209.165.200.5/0 192.168.100.2/650 100005 0:30:00

```

The entry in the LOCAL column shows the IP address of the client or server on the inside interface, while the value in the FOREIGN column shows the IP address of the client or server on the outside interface.

To view information about the Sun RPC services running on a Sun RPC server, enter the **rpcinfo -p** command from the Linux or UNIX server command line. The following is sample output from the **rpcinfo -p** command:

```

sunrpcserver:~ # rpcinfo -p
program vers proto port
100000 2 tcp 111 portmapper
100000 2 udp 111 portmapper
100024 1 udp 632 status
100024 1 tcp 635 status
100003 2 udp 2049 nfs
100003 3 udp 2049 nfs
100003 2 tcp 2049 nfs
100003 3 tcp 2049 nfs
100021 1 udp 32771 nlockmgr
100021 3 udp 32771 nlockmgr
100021 4 udp 32771 nlockmgr
100021 1 tcp 32852 nlockmgr
100021 3 tcp 32852 nlockmgr
100021 4 tcp 32852 nlockmgr
100005 1 udp 647 mountd
100005 1 tcp 650 mountd
100005 2 udp 647 mountd
100005 2 tcp 650 mountd
100005 3 udp 647 mountd
100005 3 tcp 650 mountd

```

In this output, port 647 corresponds to the mountd daemon running over UDP. The mountd process would more commonly be using port 32780. The mountd process running over TCP uses port 650 in this example.

## TFTP Inspection

TFTP inspection is enabled by default.

TFTP, described in RFC 1350, is a simple protocol to read and write files between a TFTP server and client.

The security appliance inspects TFTP traffic and dynamically creates connections and translations, if necessary, to permit file transfer between a TFTP client and server. Specifically, the inspection engine inspects TFTP read request (RRQ), write request (WRQ), and error notification (ERROR).

A dynamic secondary channel and a PAT translation, if necessary, are allocated on a reception of a valid read (RRQ) or write (WRQ) request. This secondary channel is subsequently used by TFTP for file transfer or error notification.

Only the TFTP server can initiate traffic over the secondary channel, and at most one incomplete secondary channel can exist between the TFTP client and server. An error notification from the server closes the secondary channel.

TFTP inspection must be enabled if static PAT is used to redirect TFTP traffic.

# XDMCP Inspection

XDMCP inspection is enabled by default; however, the XDMCP inspection engine is dependent upon proper configuration of the **established** command.

XDMCP is a protocol that uses UDP port 177 to negotiate X sessions, which use TCP when established.

For successful negotiation and start of an XWindows session, the security appliance must allow the TCP back connection from the Xhosted computer. To permit the back connection, use the **established** command on the security appliance. Once XDMCP negotiates the port to send the display, The **established** command is consulted to verify if this back connection should be permitted.

During the XWindows session, the manager talks to the display Xserver on the well-known port 6000 |  $n$ . Each display has a separate connection to the Xserver, as a result of the following terminal setting.

```
setenv DISPLAY Xserver: n
```

where  $n$  is the display number.

When XDMCP is used, the display is negotiated using IP addresses, which the security appliance can NAT if needed. XDCMP inspection does not support PAT.



## CHAPTER 26

# Configuring ARP Inspection and Bridging Parameters

---

### Transparent Firewall Mode Only

This chapter describes how to enable ARP inspection and how to customize bridging operations for the security appliance. In multiple context mode, the commands in this chapter can be entered in a security context, but not the system.

This chapter includes the following sections:

- [Configuring ARP Inspection, page 26-1](#)
- [Customizing the MAC Address Table, page 26-3](#)

## Configuring ARP Inspection

This section describes ARP inspection and how to enable it, and includes the following topics:

- [ARP Inspection Overview, page 26-1](#)
- [Adding a Static ARP Entry, page 26-2](#)
- [Enabling ARP Inspection, page 26-2](#)

## ARP Inspection Overview

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

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

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

**Note**

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

ARP inspection prevents malicious users from impersonating other hosts or routers (known as ARP spoofing). ARP spoofing can enable a “man-in-the-middle” attack. For example, a host sends an ARP request to the gateway router; the gateway router responds with the gateway router MAC address. The attacker, however, sends another ARP response to the host with the attacker MAC address instead of the router MAC address. The attacker can now intercept all the host traffic before forwarding it on to the router.

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

## Adding a Static ARP Entry

ARP inspection compares ARP packets with static ARP entries in the ARP table. Although hosts identify a packet destination by an IP address, the actual delivery of the packet on Ethernet relies on the Ethernet MAC address. When a router or host wants to deliver a packet on a directly connected network, it sends an ARP request asking for the MAC address associated with the IP address, and then delivers the packet to the MAC address according to the ARP response. The host or router keeps an ARP table so it does not have to send ARP requests for every packet it needs to deliver. The ARP table is dynamically updated whenever ARP responses are sent on the network, and if an entry is not used for a period of time, it times out. If an entry is incorrect (for example, the MAC address changes for a given IP address), the entry times out before it can be updated.

**Note**

The transparent firewall uses dynamic ARP entries in the ARP table for traffic to and from the security appliance, such as management traffic.

To add a static ARP entry, enter the following command:

```
hostname(config)# arp interface_name ip_address mac_address
```

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

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

## Enabling ARP Inspection

To enable ARP inspection, enter the following command:

```
hostname(config)# arp-inspection interface_name enable [flood | no-flood]
```

Where **flood** forwards non-matching ARP packets out all interfaces, and **no-flood** drops non-matching packets.

**Note**

The default setting is to flood non-matching packets. To restrict ARP through the security appliance to only static entries, then set this command to **no-flood**.

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

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

To view the current settings for ARP inspection on all interfaces, enter the **show arp-inspection** command.

## Customizing the MAC Address Table

This section describes the MAC address table, and includes the following topics:

- [MAC Address Table Overview, page 26-3](#)
- [Adding a Static MAC Address, page 26-3](#)
- [Setting the MAC Address Timeout, page 26-4](#)
- [Disabling MAC Address Learning, page 26-4](#)
- [Viewing the MAC Address Table, page 26-4](#)

## MAC Address Table Overview

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

The ASA 5505 adaptive security appliance includes a built-in switch; the switch MAC address table maintains the MAC address-to-switch port mapping for traffic within each VLAN. This section discusses the bridge MAC address table, which maintains the MAC address-to-VLAN interface mapping for traffic that passes between VLANs.

Because the security appliance is a firewall, if the destination MAC address of a packet is not in the table, the security appliance does not flood the original packet on all interfaces as a normal bridge does. Instead, it generates the following packets for directly connected devices or for remote devices:

- Packets for directly connected devices—The security appliance generates an ARP request for the destination IP address, so that the security appliance can learn which interface receives the ARP response.
- Packets for remote devices—The security appliance generates a ping to the destination IP address so that the security appliance can learn which interface receives the ping reply.

The original packet is dropped.

## Adding a Static MAC Address

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

then the security appliance drops the traffic and generates a system message. When you add a static ARP entry (see the [“Adding a Static ARP Entry” section on page 26-2](#)), a static MAC address entry is automatically added to the MAC address table.

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

```
hostname(config)# mac-address-table static interface_name mac_address
```

The *interface\_name* is the source interface.

## Setting the MAC Address Timeout

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

```
hostname(config)# mac-address-table aging-time timeout_value
```

The *timeout\_value* (in minutes) is between 5 and 720 (12 hours). 5 minutes is the default.

## Disabling MAC Address Learning

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

To disable MAC address learning, enter the following command:

```
hostname(config)# mac-learn interface_name disable
```

The **no** form of this command reenables MAC address learning. The **clear configure mac-learn** command reenables MAC address learning on all interfaces.

## Viewing the MAC Address Table

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

```
hostname# show mac-address-table [interface_name]
```

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

```
hostname# show mac-address-table
interface mac address type Time Left

outside 0009.7cbe.2100 static -
inside 0010.7cbe.6101 static -
inside 0009.7cbe.5101 dynamic 10
```

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

```
hostname# show mac-address-table inside
interface mac address type Time Left

inside 0010.7cbe.6101 static -
```

```
inside 0009.7cbe.5101 dynamic 10
```







## **PART 3**

### **Configuring VPN**





## CHAPTER 27

# Configuring IPsec and ISAKMP

---

This chapter describes how to configure the IPsec and ISAKMP standards to build Virtual Private Networks. It includes the following sections:

- [Tunneling Overview, page 27-1](#)
- [IPsec Overview, page 27-2](#)
- [Configuring ISAKMP, page 27-2](#)
- [Configuring Certificate Group Matching, page 27-9](#)
- [Configuring IPsec, page 27-11](#)
- [Clearing Security Associations, page 27-27](#)
- [Clearing Crypto Map Configurations, page 27-27](#)
- [Supporting the Nokia VPN Client, page 27-28](#)

## Tunneling Overview

Tunneling makes it possible to use a public TCP/IP network, such as the Internet, to create secure connections between remote users and a private corporate network. Each secure connection is called a tunnel.

The security appliance uses the ISAKMP and IPsec tunneling standards to build and manage tunnels. ISAKMP and IPsec accomplish the following:

- Negotiate tunnel parameters
- Establish tunnels
- Authenticate users and data
- Manage security keys
- Encrypt and decrypt data
- Manage data transfer across the tunnel
- Manage data transfer inbound and outbound as a tunnel endpoint or router

The security appliance functions as a bidirectional tunnel endpoint. It can receive plain packets from the private network, encapsulate them, create a tunnel, 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 from the public network, unencapsulate them, and send them to their final destination on the private network.

# IPsec Overview

The security appliance uses IPsec for LAN-to-LAN VPN connections, and provides the option of using IPsec for client-to-LAN VPN connections. In IPsec terminology, a *peer* is a remote-access client or another secure gateway. For both connection types, the security appliance supports only Cisco peers. Because we adhere to VPN industry standards, ASAs may work with other vendors' peers; however, we do not support them.

During tunnel establishment, the two peers negotiate security associations that govern authentication, encryption, encapsulation, and key management. These negotiations involve two phases: first, to establish the tunnel (the IKE SA); and second, to govern traffic within the tunnel (the IPsec SA).

A LAN-to-LAN VPN connects networks in different geographic locations. In IPsec LAN-to-LAN connections, the security appliance can function as initiator or responder. In IPsec client-to-LAN connections, the security appliance functions only as responder. Initiators propose SAs; responders accept, reject, or make counter-proposals—all in accordance with configured SA parameters. To establish a connection, both entities must agree on the SAs.

## Configuring ISAKMP

This section describes the Internet Key Exchange protocol which is also called the Internet Security Association and Key Management Protocol. The security appliance IKE commands use ISAKMP as a keyword, which this guide echoes. ISAKMP works with IPsec to make VPNs more scalable. This section includes the following topics:

- [ISAKMP Overview, page 27-2](#)
- [Configuring ISAKMP Policies, page 27-5](#)
- [Enabling ISAKMP on the Outside Interface, page 27-6](#)
- [Disabling ISAKMP in Aggressive Mode, page 27-6](#)
- [Determining an ID Method for ISAKMP Peers, page 27-6](#)
- [Enabling IPsec over NAT-T, page 27-7](#)
- [Enabling IPsec over TCP, page 27-8](#)
- [Waiting for Active Sessions to Terminate Before Rebooting, page 27-9](#)
- [Alerting Peers Before Disconnecting, page 27-9](#)

## ISAKMP Overview

IKE, also called ISAKMP, is the negotiation protocol that lets two hosts agree on how to build an IPsec security association. ISAKMP separates negotiation into two phases: Phase 1 and Phase 2.

Phase 1 creates the first tunnel, which protects later ISAKMP negotiation messages. Phase 2 creates the tunnel that protects data.

To set the terms of the ISAKMP negotiations, you create an ISAKMP policy, which includes the following:

- An authentication method, to ensure the identity of the peers.
- An encryption method, to protect the data and ensure privacy.

- A Hashed Message Authentication Codes (HMAC) method to ensure the identity of the sender, and to ensure that the message has not been modified in transit.
- A Diffie-Hellman group to determine the strength of the encryption-key-determination algorithm. The security appliance uses this algorithm to derive the encryption and hash keys.
- A limit to the time the security appliance uses an encryption key before replacing it.

Table 27-1 provides information about the ISAKMP policy keywords and their values.

**Table 27-1 ISAKMP Policy Keywords for CLI Commands**

| Command                                    | Keyword                   | Meaning                                                                   | Description                                                                                                                                                                                                                         |
|--------------------------------------------|---------------------------|---------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>crypto isakmp policy authentication</b> | rsa-sig                   | A digital certificate with keys generated by the RSA signatures algorithm | Specifies the authentication method the security appliance uses to establish the identity of each IPsec peer.                                                                                                                       |
|                                            | crack                     | Challenge/Response for Authenticated Cryptographic Keys                   | CRACK provides strong mutual authentication when the client authenticates using a legacy method such as RADIUS and the server uses public key authentication.                                                                       |
|                                            | pre-share<br>(default)    | Preshared keys                                                            | Preshared keys do not scale well with a growing network but are easier to set up in a small network.                                                                                                                                |
| <b>crypto isakmp policy encryption</b>     | des                       | 56-bit DES-CBC                                                            | Specifies the symmetric encryption algorithm that protects data transmitted between two IPsec peers. The default is 168-bit Triple DES.                                                                                             |
|                                            | 3des (default)            | 168-bit Triple DES                                                        |                                                                                                                                                                                                                                     |
|                                            | aes<br>aes-192<br>aes-256 |                                                                           | The Advanced Encryption Standard supports key lengths of 128, 192, 256 bits.                                                                                                                                                        |
| <b>crypto isakmp policy hash</b>           | sha (default)             | SHA-1 (HMAC variant)                                                      | Specifies the hash algorithm used to ensure data integrity. It ensures that a packet comes from where it says it comes from, and that it has not been modified in transit.                                                          |
|                                            | md5                       | MD5 (HMAC variant)                                                        | The default is SHA-1. MD5 has a smaller digest and is considered to be slightly faster than SHA-1. A successful (but extremely difficult) attack against MD5 has occurred; however, the HMAC variant IKE uses prevents this attack. |

**Table 27-1** *ISAKMP Policy Keywords for CLI Commands (continued)*

| Command                              | Keyword                            | Meaning                                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------|------------------------------------|----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>crypto isakmp policy group</b>    | 1                                  | Group 1 (768-bit)                                  | Specifies the Diffie-Hellman group identifier, which the two IPsec peers use to derive a shared secret without transmitting it to each other.<br><br>With the exception of Group 7, the lower the Diffie-Hellman group no., the less CPU time it requires to execute. The higher the Diffie-Hellman group no., the greater the security.<br><br>Cisco VPN Client Version 3.x or higher requires a minimum of Group 2. (If you configure DH Group 1, the Cisco VPN Client cannot connect.)<br><br>AES support is available on security appliances licensed for VPN-3DES only. To support the large key sizes required by AES, ISAKMP negotiation should use Diffie-Hellman (DH) Group 5.<br><br>Designed for devices with low processing power, such as PDAs and mobile telephones, Group 7 provides the greatest security. The Certicom Movian Client requires Group 7. |
|                                      | 2 (default)                        | Group 2 (1024-bit)                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|                                      | 5                                  | Group 5 (1536-bit)                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|                                      | 7                                  | Group 7 (Elliptical curve field size is 163 bits.) |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>crypto isakmp policy lifetime</b> | integer value<br>(86400 = default) | 120 to 2147483647 seconds                          | Specifies the SA lifetime. The default is 86,400 seconds or 24 hours. As a general rule, a shorter lifetime provides more secure ISAKMP negotiations (up to a point). However, with shorter lifetimes, the security appliance sets up future IPsec SAs more quickly.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Each configuration supports a maximum of 20 ISAKMP policies, each with a different set of values. Assign a unique priority to each policy you create. The lower the priority number, the higher the priority.

When ISAKMP negotiations begin, the peer that initiates the negotiation sends all of its policies to the remote peer, and the remote peer tries to find a match. The remote peer checks all of the peer's policies against each of its configured policies in priority order (highest priority first) until it discovers a match.

A match exists when both policies from the two peers contain the same encryption, hash, authentication, and Diffie-Hellman parameter values, and when the remote peer policy specifies a lifetime less than or equal to the lifetime in the policy the initiator sent. If the lifetimes are not identical, the security appliance uses the shorter lifetime. If no acceptable match exists, ISAKMP refuses negotiation and the SA is not established.

There is an implicit trade-off between security and performance when you choose a specific value for each parameter. The level of security the default values provide is adequate for the security requirements of most organizations. If you are interoperating with a peer that supports only one of the values for a parameter, your choice is limited to that value.

**Note**

New ASA configurations do not have a default ISAKMP policy.

## Configuring ISAKMP Policies

To configure ISAKMP policies, in global configuration mode, use the **crypto isakmp policy** command with its various arguments. The syntax for ISAKMP policy commands is as follows:

**crypto isakmp policy** *priority* *attribute\_name* [*attribute\_value* | *integer*]

You must include the priority in each of the ISAKMP commands. The priority number uniquely identifies the policy, and determines the priority of the policy in ISAKMP negotiations.

To enable and configure ISAKMP, complete the following steps, using the examples as a guide:

**Note**

If you do not specify a value for a given policy parameter, the default value applies.

- Step 1** Specify the encryption algorithm. The default is Triple DES. This example sets encryption to DES.
- ```
crypto isakmp policy priority encryption [aes | aes-192 | aes-256 | des | 3des]
```
- For example:
- ```
hostname(config)# crypto isakmp policy 2 encryption des
```
- Step 2** Specify the hash algorithm. The default is SHA-1. This example configures MD5.
- ```
crypto isakmp policy priority hash [md5 | sha]
```
- For example:
- ```
hostname(config)# crypto isakmp policy 2 hash md5
```
- Step 3** Specify the authentication method. The default is preshared keys. This example configures RSA signatures.
- ```
crypto isakmp policy priority authentication [pre-share | crack | rsa-sig]
```
- For example:
- ```
hostname(config)# crypto isakmp policy 2 authentication rsa-sig
```
- Step 4** Specify the Diffie-Hellman group identifier. The default is Group 2. This example configures Group 5.
- ```
crypto isakmp policy priority group [1 | 2 | 5 | 7]
```
- For example:
- ```
hostname(config)# crypto isakmp policy 2 group 5
```
- Step 5** Specify the SA lifetime. This examples sets a lifetime of 4 hours (14400 seconds). The default is 86400 seconds (24 hours).
- ```
crypto isakmp policy priority lifetime seconds
```
- For example:
- ```
hostname(config)# crypto isakmp policy 2 lifetime 14400
```

## Enabling ISAKMP on the Outside Interface

You must enable ISAKMP on the interface that terminates the VPN tunnel. Typically this is the outside, or public interface.

To enable ISAKMP, enter the following command:

```
crypto isakmp enable interface-name
```

For example:

```
hostname(config)# crypto isakmp enable outside
```

## Disabling ISAKMP in Aggressive Mode

Phase 1 ISAKMP negotiations can use either main mode or aggressive mode. Both provide the same services, but aggressive mode requires only two exchanges between the peers totaling 3 messages, rather than three exchanges totaling 6 messages. Aggressive mode is faster, but does not provide identity protection for the communicating parties. Therefore, the peers must exchange identification information prior to establishing a secure SA. Aggressive mode is enabled by default.

- Main mode is slower, using more exchanges, but it protects the identities of the communicating peers.
- Aggressive mode is faster, but does not protect the identities of the peers.

To disable ISAKMP in aggressive mode, enter the following command:

```
crypto isakmp am-disable
```

For example:

```
hostname(config)# crypto isakmp am-disable
```

If you have disabled aggressive mode, and want to revert to back to it, use the **no** form of the command.

For example:

```
hostname(config)# no crypto isakmp am-disable
```



### Note

Disabling aggressive mode prevents Cisco VPN clients from using preshared key authentication to establish tunnels to the security appliance. However, they may use certificate-based authentication (that is, ASA or RSA) to establish tunnels.

## Determining an ID Method for ISAKMP Peers

During Phase I ISAKMP negotiations the peers must identify themselves to each other. You can choose the identification method from the following options:



|                  |                                                                                                                                                                                                        |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Address</b>   | Uses the IP addresses of the hosts exchanging ISAKMP identity information.                                                                                                                             |
| <b>Automatic</b> | Determines ISAKMP negotiation by connection type: <ul style="list-style-type: none"> <li>• IP address for preshared key.</li> <li>• Cert Distinguished Name for certificate authentication.</li> </ul> |
| <b>Hostname</b>  | Uses the fully qualified domain name of the hosts exchanging ISAKMP identity information (default). This name comprises the hostname and the domain name.                                              |
| <b>Key ID</b>    | Uses the string the remote peer uses to look up the preshared key.                                                                                                                                     |

The security appliance uses the Phase I ID to send to the peer. This is true for all VPN scenarios except LAN-to-LAN connections in main mode that authenticate with preshared keys.

The default setting is hostname.

To change the peer identification method, enter the following command:

```
crypto isakmp identity {address | hostname | key-id id-string | auto}
```

For example, the following command sets the peer identification method to automatic:

```
hostname(config)# crypto isakmp identity auto
```

## Enabling IPsec over NAT-T

NAT-T lets IPsec peers establish a connection through a NAT device. It does this by encapsulating IPsec traffic in UDP datagrams, using port 4500, thereby providing NAT devices with port information. NAT-T auto-detects any NAT devices, and only encapsulates IPsec traffic when necessary. This feature is disabled by default.

With the exception of the home zone on the Cisco ASA 5505, the security appliance can simultaneously support standard IPsec, IPsec over TCP, NAT-T, and IPsec over UDP, depending on the client with which it is exchanging data. When both NAT-T and IPsec over UDP are enabled, NAT-T takes precedence. IPsec over TCP, if enabled, takes precedence over all other connection methods.

When you enable NAT-T, the security appliance automatically opens port 4500 on all IPsec enabled interfaces.

The security appliance supports multiple IPsec peers behind a single NAT/PAT device operating in one of the following networks, but not both:

- LAN-to-LAN
- Remote access

In a mixed environment, the remote access tunnels fail the negotiation because all peers appear to be coming from the same public IP address, that of the NAT device. Also, remote access tunnels fail in a mixed environment because they often use the same name as the LAN-to-LAN tunnel group (that is, the IP address of the NAT device). This match can cause negotiation failures among multiple peers in a mixed LAN-to-LAN and remote access network of peers behind the NAT device.

## Using NAT-T

To use NAT-T, you must perform the following tasks:

- 
- Step 1** Enter the following command to enable IPsec over NAT-T globally on the security appliance.

```
crypto isakmp nat-traversal natkeepalive
```

natkeepalive is in the range 10 to 3600 seconds. The default is 20 seconds.

For example, enter the following command to enable NAT-T and set the keepalive to one hour.

```
hostname(config)# crypto isakmp nat-traversal 3600
```

**Step 2** Select the “before-fragmentation” option for the IPsec fragmentation policy.

This option lets traffic travel across NAT devices that do not support IP fragmentation. It does not impede the operation of NAT devices that do support IP fragmentation.

## Enabling IPsec over TCP

IPsec over TCP enables a Cisco VPN client to operate in an environment in which standard ESP or ISAKMP cannot function, or can function only with modification to existing firewall rules. IPsec over TCP encapsulates both the ISAKMP and IPsec protocols within a TCP-like packet, and enables secure tunneling through both NAT and PAT devices and firewalls. This feature is disabled by default.



### Note

This feature does not work with proxy-based firewalls.

IPsec over TCP works with remote access clients. You enable it globally, and it works on all ISAKMP enabled interfaces. It is a client to security appliance feature only. It does not work for LAN-to-LAN connections.

The security appliance can simultaneously support standard IPsec, IPsec over TCP, NAT-Traversal, and IPsec over UDP, depending on the client with which it is exchanging data. IPsec over TCP, if enabled, takes precedence over all other connection methods.

The VPN 3002 hardware client, which supports one tunnel at a time, can connect using standard IPsec, IPsec over TCP, NAT-Traversal, or IPsec over UDP.

You enable IPsec over TCP on both the security appliance and the client to which it connects.

You can enable IPsec over TCP for up to 10 ports that you specify. If you enter a well-known port, for example port 80 (HTTP) or port 443 (HTTPS), the system displays a warning that the protocol associated with that port no longer works on the public interface. The consequence is that you can no longer use a browser to manage the security appliance through the public interface. To solve this problem, reconfigure the HTTP/HTTPS management to different ports.

The default port is 10000.

You must configure TCP port(s) on the client as well as on the security appliance. The client configuration must include at least one of the ports you set for the security appliance.

To enable IPsec over TCP globally on the security appliance, enter the following command:

```
crypto isakmp ipsec-over-tcp [port port 1...port0]
```

This example enables IPsec over TCP on port 45:

```
hostname(config)# crypto isakmp tcp port 45
```

## Waiting for Active Sessions to Terminate Before Rebooting

You can schedule a security appliance reboot to occur only when all active sessions have terminated voluntarily. This feature is disabled by default.

To enable waiting for all active sessions to voluntarily terminate before the security appliance reboots, enter the following command:

```
crypto isakmp reload-wait
```

For example:

```
hostname(config)# crypto isakmp reload-wait
```

Use the **reload** command to reboot the security appliance. If you set the **reload-wait** command, you can use the **reload quick** command to override the **reload-wait** setting. The **reload** and **reload-wait** commands are available in privileged EXEC mode; neither includes the **isakmp** prefix.

## Alerting Peers Before Disconnecting

Remote access or LAN-to-LAN sessions can drop for several reasons, such as: a security appliance shutdown or reboot, session idle timeout, maximum connection time exceeded, or administrator cut-off.

The security appliance can notify qualified peers (in LAN-to-LAN configurations), Cisco VPN clients and VPN 3002 hardware clients of sessions that are about to be disconnected. The peer or client receiving the alert decodes the reason and displays it in the event log or in a pop-up pane. This feature is disabled by default.

Qualified clients and peers include the following:

- Security appliances with Alerts enabled.
- Cisco VPN clients running version 4.0 or later software (no configuration required).
- VPN 3002 hardware clients running version 4.0 or later software, and with Alerts enabled.
- VPN 3000 series concentrators running version 4.0 or later software, with Alerts enabled.

To enable disconnect notification to IPsec peers, enter the **crypto isakmp disconnect-notify** command.

For example:

```
hostname(config)# crypto isakmp disconnect-notify
```

## Configuring Certificate Group Matching

Tunnel groups define user connection terms and permissions. Certificate group matching lets you match a user to a tunnel group using either the Subject DN or Issuer DN of the user certificate.

To match users to tunnel groups based on these fields of the certificate, you must first create rules that define a matching criteria, and then associate each rule with the desired tunnel group.

To create a certificate map, use the **crypto ca certificate map** command. To define a tunnel group, use the **tunnel-group** command.

You must also configure a certificate group matching policy that sets one of the following methods for identifying the permission groups of certificate users:

- Match the group from the rules

- Match the group from the organizational unit (OU) field
- Use a default group for all certificate users

You can use any or all of these methods.

## Creating a Certificate Group Matching Rule and Policy

To configure the policy and rules by which certificate-based ISAKMP sessions map to tunnel groups, and to associate the certificate map entries with tunnel groups, enter the **tunnel-group-map** command in global configuration mode.

The syntax follows:

**tunnel-group-map enable** { *rules* | *ou* | *ike-id* | *peer ip* }

**tunnel-group-map** [*rule-index*] **enable policy**

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>policy</i>     | <p>Specifies the policy for deriving the tunnel group name from the certificate. <i>Policy</i> can be one of the following:</p> <p><i>ike-id</i>—Indicates that if a tunnel-group is not determined based on a rule lookup or taken from the <i>ou</i>, then the certificate-based ISAKMP sessions are mapped to a tunnel group based on the content of the phase1 ISAKMP ID.</p> <p><i>ou</i>—Indicates that if a tunnel-group is not determined based on a rule lookup, then use the value of the OU in the subject distinguished name (DN).</p> <p><i>peer-ip</i>—Indicates that if a tunnel-group is not determined based on a rule lookup or taken from the <i>ou</i> or <i>ike-id</i> methods, then use the peer IP address.</p> <p><i>rules</i>—Indicates that the certificate-based ISAKMP sessions are mapped to a tunnel group based on the certificate map associations configured by this command.</p> |
| <i>rule index</i> | (Optional) Refers to parameters specified by the <b>crypto ca certificate map</b> command. The values are 1 to 65535.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Be aware of the following:

- You can invoke this command multiple times as long as each invocation is unique and you do not reference a map index more than once.
- Rules cannot be longer than 255 characters.
- You can assign multiple rules to the same group. To do that, you add the rule priority and group first. Then you define as many criteria statements as you need for each group. When multiple rules are assigned to the same group, a match results for the first rule that tests true.
- Create a single rule if you want to require all criteria to match before assigning a user to a specific tunnel group. Requiring all criteria to match is equivalent to a logical AND operation. Alternatively, create one rule for each criterion if you want to require that only one match before assigning a user to a specific tunnel group. Requiring only one criterion to match is equivalent to a logical OR operation.

The following example enables mapping of certificate-based ISAKMP sessions to a tunnel group based on the content of the phase1 ISAKMP ID:

```
hostname(config)# tunnel-group-map enable ike-id
hostname(config)#
```

The following example enables mapping of certificate-based ISAKMP sessions to a tunnel group based on the IP address of the peer:

```
hostname(config)# tunnel-group-map enable peer-ip
hostname(config)#
```

The following example enables mapping of certificate-based ISAKMP sessions based on the organizational unit (OU) in the subject distinguished name (DN):

```
hostname(config)# tunnel-group-map enable ou
hostname(config)#
```

The following example enables mapping of certificate-based ISAKMP sessions based on established rules:

```
hostname(config)# tunnel-group-map enable rules
hostname(config)#
```

## Using the Tunnel-group-map default-group Command

This command specifies a default tunnel group to use when the configuration does not specify a tunnel group.

The syntax is **tunnel-group-map** [*rule-index*] **default-group** *tunnel-group-name* where the *rule-index* is the priority for the rule, and *tunnel-group name* must be for a tunnel group that already exists.

## Configuring IPsec

This section provides background information about IPsec and describes the procedures required to configure the security appliance when using IPsec to implement a VPN. It contains the following topics:

- [Understanding IPsec Tunnels, page 27-11](#)
- [Understanding Transform Sets, page 27-12](#)
- [Defining Crypto Maps, page 27-12](#)
- [Applying Crypto Maps to Interfaces, page 27-20](#)
- [Using Interface Access Lists, page 27-20](#)
- [Changing IPsec SA Lifetimes, page 27-22](#)
- [Creating a Basic IPsec Configuration, page 27-22](#)
- [Using Dynamic Crypto Maps, page 27-24](#)
- [Providing Site-to-Site Redundancy, page 27-26](#)
- [Viewing an IPsec Configuration, page 27-26](#)

## Understanding IPsec Tunnels

IPsec tunnels are sets of SAs that the security appliance establishes between peers. The SAs define the protocols and algorithms to apply to sensitive data, and also specify the keying material the peers use. IPsec SAs control the actual transmission of user traffic. SAs are unidirectional, but are generally established in pairs (inbound and outbound).

The peers negotiate the settings to use for each SA. Each SA consists of the following:

- Transform sets
- Crypto maps
- Access lists
- Tunnel groups
- Prefragmentation policies

## Understanding Transform Sets

A transform set is a combination of security protocols and algorithms that define how the security appliance protects data. During IPsec SA negotiations, the peers must identify a transform set that is the same at both peers. The security appliance then applies the matching transform set to create an SA that protects data flows in the access list for that crypto map.

The security appliance tears down the tunnel if you change the definition of the transform set used to create its SA. See “[Clearing Security Associations](#)” for further information.



### Note

If you clear or delete the only element in a transform set, the security appliance automatically removes the crypto map references to it.

## Defining Crypto Maps

*Crypto maps* define the IPsec policy to be negotiated in the IPsec SA. They include the following:

- Access list to identify the packets that the IPsec connection permits and protects.
- Peer identification
- Local address for the IPsec traffic (See “[Applying Crypto Maps to Interfaces](#)” for more details.)
- Up to six transform sets with which to attempt to match the peer security settings.

A *crypto map set* consists of one or more crypto maps that have the same map name. You create a crypto map set when you create its first crypto map. The following command syntax creates or adds to a crypto map:

```
crypto map map-name seq-num match address access-list-name
```

You can continue to enter this command to add crypto maps to the crypto map set. In the following example, “mymap” is the name of the crypto map set to which you might want to add crypto maps:

```
crypto map mymap 10 match address 101
```

The *sequence number* (seq-num) shown in the syntax above distinguishes one crypto map from another one with the same name. The sequence number assigned to a crypto map also determines its priority among the other crypto maps within a crypto map set. The lower the sequence number, the higher the priority. After you assign a crypto map set to an interface, the security appliance evaluates all IP traffic passing through the interface against the crypto maps in the set, beginning with the crypto map with the lowest sequence number.

The ACL assigned to a crypto map consists of all of the ACEs that have the same access-list-name, as shown in the following command syntax:

```
access-list access-list-name {deny | permit} ip source source-netmask destination
destination-netmask
```

Each ACL consists of one or more ACEs that have the same access-list-name. You create an ACL when you create its first ACE. The following command syntax creates or adds to an ACL:

```
access-list access-list-name {deny | permit} ip source source-netmask destination destination-netmask
```

In the following example, the security appliance applies the IPsec protections assigned to the crypto map to all traffic flowing from the 10.0.0.0 subnet to the 10.1.1.0 subnet.

```
access-list 101 permit ip 10.0.0.0 255.255.255.0 10.1.1.0 255.255.255.0
```

The crypto map that matches the packet determines the security settings used in the SA negotiations. If the local security appliance initiates the negotiation, it uses the policy specified in the static crypto map to create the offer to send to the specified peer. If the peer initiates the negotiation, the security appliance attempts to match the policy to a static crypto map, and if that fails, any dynamic crypto maps in the crypto map set, to decide whether to accept or reject the peer offer.

For two peers to succeed in establishing an SA, they must have at least one compatible crypto map. To be compatible, a crypto map must meet the following criteria:

- The crypto map must contain compatible crypto ACLs (for example, mirror image ACLs). If the responding peer uses dynamic crypto maps, so must the security appliance as a requirement to apply IPsec.
- Each crypto map identifies the other peer (unless the responding peer uses dynamic crypto maps).
- The crypto maps have at least one transform set in common.

You can apply only one crypto map set to a single interface. Create more than one crypto map for a particular interface on the security appliance if any of the following conditions exist:

- You want specific peers to handle different data flows.
- You want different IPsec security to apply to different types of traffic.

For example, create a crypto map and assign an ACL to identify traffic between two subnets and assign one transform set. Create another crypto map with a different ACL to identify traffic between another two subnets and apply a transform set with different VPN parameters.

If you create more than one crypto map for an interface, specify a sequence number (seq-num) for each map entry to determine its priority within the crypto map set.

Each ACE contains a permit or deny statement. [Table 27-2](#) explains the special meanings of permit and deny ACEs in ACLs applied to crypto maps.

**Table 27-2**      *Special Meanings of Permit and Deny in Crypto Access Lists Applied to Outbound Traffic*

| Result of Crypto Map Evaluation                            | Response                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Match criterion in an ACE containing a permit statement    | Halt further evaluation of the packet against the remaining ACEs in the crypto map set, and evaluate the packet security settings against those in the transform sets assigned to the crypto map. After matching the security settings to those in a transform set, the security appliance applies the associated IPsec settings. Typically for outbound traffic, this means that it decrypts, authenticates, and routes the packet. |
| Match criterion in an ACE containing a deny statement      | Interrupt further evaluation of the packet against the remaining ACEs in the crypto map under evaluation, and resume evaluation against the ACEs in the next crypto map, as determined by the next seq-num assigned to it.                                                                                                                                                                                                           |
| Fail to match all tested permit ACEs in the crypto map set | Route the packet without encrypting it.                                                                                                                                                                                                                                                                                                                                                                                              |

ACEs containing deny statements filter out outbound traffic that does not require IPsec protection (for example, routing protocol traffic). Therefore, insert initial deny statements to filter outbound traffic that should not be evaluated against permit statements in a crypto access list.

For an inbound, encrypted packet, the security appliance uses the source address and ESP SPI to determine the decryption parameters. After the security appliance decrypts the packet, it compares the inner header of the decrypted packet to the permit ACEs in the ACL associated with the packet SA. If the inner header fails to match the proxy, the security appliance drops the packet. If the inner header matches the proxy, the security appliance routes the packet.


When comparing the inner header of an inbound packet that was not encrypted, the security appliance ignores all deny rules because they would prevent the establishment of a Phase 2 SA.

**Note**

To route inbound, unencrypted traffic as clear text, insert deny ACEs before permit ACEs.

Figure 27-1 shows an example LAN-to-LAN network of security appliances.



**Figure 27-1** Effect of Permit and Deny ACEs on Traffic (Conceptual Addresses)

143514

The simple address notation shown in this figure and used in the following explanation is an abstraction. An example with real IP addresses follows the explanation.

The objective in configuring Security Appliances A, B, and C in this example LAN-to-LAN network is to permit tunneling of all traffic originating from one of the hosts shown in Figure 27-1 and destined for one of the other hosts. However, because traffic from Host A.3 contains sensitive data from the Human Resources department, it requires strong encryption and more frequent rekeying than the other traffic. So we want to assign a special transform set for traffic from Host A.3.

To configure Security Appliance A for outbound traffic, we create two crypto maps, one for traffic from Host A.3 and the other for traffic from the other hosts in Network A, as shown in the following example:


```
Crypto Map Seq_No_1
 deny packets from A.3 to B
 deny packets from A.3 to C
 permit packets from A to B
 permit packets from A to C
Crypto Map Seq_No_2
 permit packets from A.3 to B
 permit packets from A.3 to C
```

After creating the ACLs, you assign a transform set to each crypto map to apply the required IPsec to each matching packet.

Cascading ACLs involves the insertion of deny ACEs to bypass evaluation against an ACL and resume evaluation against a subsequent ACL in the crypto map set. Because you can associate each crypto map with different IPsec settings, you can use deny ACEs to exclude special traffic from further evaluation in the corresponding crypto map, and match the special traffic to permit statements in another crypto map to provide or require different security. The sequence number assigned to the crypto ACL determines its position in the evaluation sequence within the crypto map set.

Figure 27-2 shows the cascading ACLs created from the conceptual ACEs above. The meaning of each symbol in the figure follows.

|                                                                                   |                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Crypto map within a crypto map set.                                                                                                                                                                                                     |
|  | (Gap in a straight line) Exit from a crypto map when a packet matches an ACE.                                                                                                                                                           |
|  | Packet that fits the description of one ACE. Each size ball represents a different packet matching the respective ACE in the figure. The differences in size merely represent differences in the source and destination of each packet. |
|  | Redirection to the next crypto map in the crypto map set.                                                                                                                                                                               |
|  | Response when a packet either matches an ACE or fails to match all of the permit ACEs in a crypto map set.                                                                                                                              |

**Figure 27-2** Cascading ACLs in a Crypto Map Set

Security Appliance A evaluates a packet originating from Host A.3 until it matches a permit ACE and attempts to assign the IPsec security associated with the crypto map. Whenever the packet matches a deny ACE, the security appliance ignores the remaining ACEs in the crypto map and resumes evaluation against the next crypto map, as determined by the sequence number assigned to it. So in the example, if Security Appliance A receives a packet from Host A.3, it matches the packet to a deny ACE in the first crypto map and resumes evaluation of the packet against the next crypto map. When it matches the packet to the permit ACE in that crypto map, it applies the associated IPsec security (strong encryption and frequent rekeying).


To complete the security appliance configuration in the example network, we assign mirror crypto maps to Security Appliances B and C. However, because security appliances ignore deny ACEs when evaluating inbound, encrypted traffic, we can omit the mirror equivalents of the deny A.3 B and deny A.3 C ACEs, and therefore omit the mirror equivalents of Crypto Map 2. So the configuration of cascading ACLs in Security Appliances B and C is unnecessary.

Table 27-3 shows the ACLs assigned to the crypto maps configured for all three security appliances in Figure 27-1.

**Table 27-3** Example Permit and Deny Statements (Conceptual)

| Security Appliance A    |              | Security Appliance B    |             | Security Appliance C    |             |
|-------------------------|--------------|-------------------------|-------------|-------------------------|-------------|
| Crypto Map Sequence No. | ACE Pattern  | Crypto Map Sequence No. | ACE Pattern | Crypto Map Sequence No. | ACE Pattern |
| 1                       | deny A.3 B   | 1                       | permit B A  | 1                       | permit C A  |
|                         | deny A.3 C   |                         |             |                         |             |
|                         | permit A B   |                         |             |                         |             |
|                         | permit A C   |                         | permit B C  |                         | permit C B  |
| 2                       | permit A.3 B |                         |             |                         |             |
|                         | permit A.3 C |                         |             |                         |             |

Figure 27-3 maps the conceptual addresses shown in Figure 27-1 to real IP addresses.

**Figure 27-3** Effect of Permit and Deny ACEs on Traffic (Real Addresses)

The tables that follow combine the IP addresses shown in [Figure 27-3](#) to the concepts shown in [Table 27-3](#). The real ACEs shown in these tables ensure that all IPsec packets under evaluation within this network receive the proper IPsec settings.

**Table 27-4** Example Permit and Deny Statements for Security Appliance A

| Security Appliance | Crypto Map Sequence No. | ACE Pattern  | Real ACEs                                                         |
|--------------------|-------------------------|--------------|-------------------------------------------------------------------|
| A                  | 1                       | deny A.3 B   | deny 192.168.3.3 255.255.255.192 192.168.12.0 255.255.255.248     |
|                    |                         | deny A.3 C   | deny 192.168.3.3 255.255.255.192 192.168.201.0 255.255.255.224    |
|                    |                         | permit A B   | permit 192.168.3.0 255.255.255.192 192.168.12.0 255.255.255.248   |
|                    |                         | permit A C   | permit 192.168.3.0 255.255.255.192 192.168.201.0 255.255.255.224  |
|                    | 2                       | permit A.3 B | permit 192.168.3.3 255.255.255.192 192.168.12.0 255.255.255.248   |
|                    |                         | permit A.3 C | permit 192.168.3.3 255.255.255.192 192.168.201.0 255.255.255.224  |
| B                  | None needed             | permit B A   | permit 192.168.12.0 255.255.255.248 192.168.3.0 255.255.255.192   |
|                    |                         | permit B C   | permit 192.168.12.0 255.255.255.248 192.168.201.0 255.255.255.224 |
| C                  | None needed             | permit C A   | permit 192.168.201.0 255.255.255.224 192.168.3.0 255.255.255.192  |
|                    |                         | permit C B   | permit 192.168.201.0 255.255.255.224 192.168.12.0 255.255.255.248 |

You can apply the same reasoning shown in the example network to use cascading ACLs to assign different security settings to different hosts or subnets protected by a Cisco security appliance.

**Note**

By default, the security appliance does not support IPsec traffic destined for the same interface from which it enters. (Names for this type of traffic include U-turn, hub-and-spoke, and hairpinning.) However, you might want IPsec to support U-turn traffic. To do so, insert an ACE to permit traffic to and from the network. For example, to support U-turn traffic on Security Appliance B, add a conceptual “permit B B” ACE to ACL1. The actual ACE would be as follows:

```
permit 192.168.12.0 255.255.255.248 192.168.12.0 255.255.255.248
```

## Applying Crypto Maps to Interfaces

You must assign a crypto map set to each interface through which IPsec traffic flows. The security appliance supports IPsec on all interfaces. Assigning the crypto map set to an interface instructs the security appliance to evaluate all the traffic against the crypto map set and to use the specified policy during connection or SA negotiation.

Assigning a crypto map to an interface also initializes run-time data structures, such as the SA database and the security policy database. Reassigning a modified crypto map to the interface resynchronizes the run-time data structures with the crypto map configuration. Also, adding new peers through the use of new sequence numbers and reassigning the crypto map does not tear down existing connections.

## Using Interface Access Lists

By default, the security appliance lets IPsec packets bypass interface ACLs. If you want to apply interface access lists to IPsec traffic, use the **no** form of the **sysopt connection permit-ipsec** command.

The crypto map access list bound to the outgoing interface either permits or denies IPsec packets through the VPN tunnel. IPsec authenticates and deciphers packets that arrive from an IPsec tunnel, and subjects them to evaluation against the ACL associated with the tunnel.

Access lists define which IP traffic to protect. For example, you can create access lists to protect all IP traffic between two subnets or two hosts. (These access lists are similar to access lists used with the **access-group** command. However, with the **access-group** command, the access list determines which traffic to forward or block at an interface.)

Before the assignment to crypto maps, the access lists are not specific to IPsec. Each crypto map references the access lists and determines the IPsec properties to apply to a packet if it matches a permit in one of the access lists.

Access lists assigned to IPsec crypto maps have four primary functions:

- Select outbound traffic to be protected by IPsec (permit = protect).
- Trigger an ISAKMP negotiation for data travelling without an established SA.
- Process inbound traffic to filter out and discard traffic that should have been protected by IPsec.
- Determine whether to accept requests for IPsec SAs when processing IKE negotiation from the peer. (Negotiation applies only to **ipsec-isakmp crypto map** entries.) The peer must “permit” a data flow associated with an **ipsec-isakmp crypto map** command entry to ensure acceptance during negotiation.

Regardless of whether the traffic is inbound or outbound, the security appliance evaluates traffic against the access lists assigned to an interface. You assign IPsec to an interface as follows:


---

**Step 1** Create the access lists to be used for IPsec.

- Step 2** Map the lists to one or more crypto maps, using the same crypto map name.
- Step 3** Map the transform sets to the crypto maps to apply IPsec to the data flows.
- Step 4** Apply the crypto maps collectively as a “crypto map set” by assigning the crypto map name they share to the interface.

In [Figure 27-4](#), IPsec protection applies to traffic between Host 10.0.0.1 and Host 10.2.2.2 as the data exits the outside interface on Security Appliance A toward Host 10.2.2.2.

**Figure 27-4** How Crypto Access Lists Apply to IPsec



Security Appliance A evaluates traffic from Host 10.0.0.1 to Host 10.2.2.2, as follows:

- source = host 10.0.0.1
- dest = host 10.2.2.2

Security Appliance A also evaluates traffic from Host 10.2.2.2 to Host 10.0.0.1, as follows:

- source = host 10.2.2.2
- dest = host 10.0.0.1

The first permit statement that matches the packet under evaluation determines the scope of the IPsec SA.



**Note**

If you delete the only element in an access list, the security appliance also removes the associated crypto map.

If you modify an access list currently referenced by one or more crypto maps, use the **crypto map interface** command to reinitialize the run-time SA database. See the **crypto map** command for more information.

We recommend that for every crypto access list specified for a static crypto map that you define at the local peer, you define a “mirror image” crypto access list at the remote peer. The crypto maps should also support common transforms and refer to the other system as a peer. This ensures correct processing of IPsec by both peers.

**Note**

Every static crypto map must define an access list and an IPsec peer. If either is missing, the crypto map is incomplete and the security appliance drops any traffic that it has not already matched to an earlier, complete crypto map. Use the **show conf** command to ensure that every crypto map is complete. To fix an incomplete crypto map, remove the crypto map, add the missing entries, and reapply it.

We discourage the use of the **any** keyword to specify source or destination addresses in crypto access lists because they cause problems. We strongly discourage the **permit any any** command statement because it does the following:

- Protects all outbound traffic, including all protected traffic sent to the peer specified in the corresponding crypto map.
- Requires protection for all inbound traffic.

In this scenario, the security appliance silently drops all inbound packets that lack IPsec protection.

Be sure that you define which packets to protect. If you use the **any** keyword in a **permit** statement, preface it with a series of **deny** statements to filter out traffic that would otherwise fall within that **permit** statement that you do not want to protect.

## Changing IPsec SA Lifetimes

You can change the global lifetime values that the security appliance uses when negotiating new IPsec SAs. You can override these global lifetime values for a particular crypto map.

IPsec SAs use a derived, shared, secret key. The key is an integral part of the SA; they time out together to require the key to refresh. Each SA has two lifetimes: “timed” and “traffic-volume.” An SA expires after the respective lifetime and negotiations begin for a new one. The default lifetimes are 28,800 seconds (eight hours) and 4,608,000 kilobytes (10 megabytes per second for one hour).

If you change a global lifetime, the security appliance drops the tunnel. It uses the new value in the negotiation of subsequently established SAs.

When a crypto map does not have configured lifetime values and the security appliance requests a new SA, it inserts the global lifetime values used in the existing SA into the request sent to the peer. When a peer receives a negotiation request, it uses the smaller of either the lifetime value the peer proposes or the locally configured lifetime value as the lifetime of the new SA.

The peers negotiate a new SA before crossing the lifetime threshold of the existing SA to ensure that a new SA is ready when the existing one expires. The peers negotiate a new SA when about 5 to 15 percent of the lifetime of the existing SA remains.

## Creating a Basic IPsec Configuration

You can create basic IPsec configurations with static or dynamic crypto maps.

To create a basic IPsec configuration using a static crypto map, perform the following steps:

**Step 1**

To create an access list to define the traffic to protect, enter the following command:

```
access-list access-list-name {deny | permit} ip source source-netmask destination destination-netmask
```

For example:



```
access-list 101 permit ip 10.0.0.0 255.255.255.0 10.1.1.0 255.255.255.0
```

In this example, the **permit** keyword causes all traffic that matches the specified conditions to be protected by crypto.

- Step 2** To configure a transform set that defines how to protect the traffic, enter the following command:

```
crypto ipsec transform-set transform-set-name transform1 [ttransform2, transform3]
```

For example:

```
crypto ipsec transform-set myset1 esp-des esp-sha-hmac
crypto ipsec transform-set myset2 esp-3des esp-sha-hmac
crypto ipsec transform-set aes_set esp-md5-hmac esp-aes-256
```

In this example, “myset1” and “myset2” and “aes\_set” are the names of the transform sets.

- Step 3** To create a crypto map, perform the following steps:

- a. Assign an access list to a crypto map:

```
crypto map map-name seq-num match address access-list-name
```

In the following example, “mymap” is the name of the crypto map set. The map set sequence number 10, which is used to rank multiple entries within one crypto map set. The lower the sequence number, the higher the priority.

```
crypto map mymap 10 match address 101
```

In this example, the access list named 101 is assigned to crypto map “mymap.”

- b. Specify the peer to which the IPsec protected traffic can be forwarded:

```
crypto map map-name seq-num set peer ip-address
```

For example:

```
crypto map mymap 10 set peer 192.168.1.100
```

The security appliance sets up an SA with the peer assigned the IP address 192.168.1.100. Specify multiple peers by repeating this command.

- c. Specify which transform sets are allowed for this crypto map. List multiple transform sets in order of priority (highest priority first). You can specify up to 11 transform sets in a crypto map.

```
crypto map map-name seq-num set transform-set transform-set-name1
[transform-set-name2, ...transform-set-name6]
```

For example:

```
crypto map mymap 10 set transform-set myset1 myset2
```

In this example, when traffic matches access list 101, the SA can use either “myset1” (first priority) or “myset2” (second priority) depending on which transform set matches the transform set of the peer.

- d. (Optional) Specify an SA lifetime for the crypto map if you want to override the global lifetime.

```
crypto map map-name seq-num set security-association lifetime {seconds seconds |
kilobytes kilobytes}
```

For example:

```
crypto map mymap 10 set security-association lifetime seconds 2700
```

This example shortens the timed lifetime for the crypto map “mymap 10” to 2700 seconds (45 minutes). The traffic volume lifetime is not changed.

- e. (Optional) Specify that IPsec require perfect forward secrecy when requesting new SA for this crypto map, or require PFS in requests received from the peer:

```
crypto map map-name seq-num set pfs [group1 | group2 | group5 | group7]
```

For example:

```
crypto map mymap 10 set pfs group2
```

This example requires PFS when negotiating a new SA for the crypto map “mymap 10.” The security appliance uses the 1024-bit Diffie-Hellman prime modulus group in the new SA.

- Step 4** Apply a crypto map set to an interface for evaluating IPsec traffic:

```
crypto map map-name interface interface-name
```

For example:

```
crypto map mymap interface outside
```

In this example, the security appliance evaluates the traffic going through the outside interface against the crypto map “mymap” to determine whether it needs to be protected.

## Using Dynamic Crypto Maps

A dynamic crypto map is a crypto map without all of the parameters configured. It acts as a policy template where the missing parameters are later dynamically learned, as the result of an IPsec negotiation, to match the peer requirements. The security appliance applies a dynamic crypto map to let a peer negotiate a tunnel if its IP address is not already identified in a static crypto map. This occurs with the following types of peers:

- Peers with dynamically assigned public IP addresses.  
Both LAN-to-LAN and remote access peers can use DHCP to obtain a public IP address. The security appliance uses this address only to initiate the tunnel.
- Peers with dynamically assigned private IP addresses.  
Peers requesting remote access tunnels typically have private IP addresses assigned by the headend. Generally, LAN-to-LAN tunnels have a predetermined set of private networks that are used to configure static maps and therefore used to establish IPsec SAs.

As an administrator configuring static crypto maps, you might not know the IP addresses that are dynamically assigned (via DHCP or some other method), and you might not know the private IP addresses of other clients, regardless of how they were assigned. VPN clients typically do not have static IP addresses; they require a dynamic crypto map to allow IPsec negotiation to occur. For example, the headend assigns the IP address to a Cisco VPN client during IKE negotiation, which the client then uses to negotiate IPsec SAs.



### Note

A dynamic crypto map requires only the **transform-set** parameter.

Dynamic crypto maps can ease IPsec configuration and we recommend them for use in networks where the peers are not always predetermined. Use dynamic crypto maps for Cisco VPN clients (such as mobile users) and routers that obtain dynamically assigned IP addresses.

**Tip**

Use care when using the **any** keyword in **permit** entries in dynamic crypto maps. If the traffic covered by such a **permit** entry could include multicast or broadcast traffic, insert **deny** entries for the appropriate address range into the access list. Remember to insert **deny** entries for network and subnet broadcast traffic, and for any other traffic that IPsec should not protect.

Dynamic crypto maps work only to negotiate SAs with remote peers that initiate the connection. The security appliance cannot use dynamic crypto maps to initiate connections to a remote peer. With a dynamic crypto map, if outbound traffic matches a permit entry in an access list and the corresponding SA does not yet exist, the security appliance drops the traffic.

A crypto map set may include a dynamic crypto map. Dynamic crypto map sets should be the lowest priority crypto maps in the crypto map set (that is, they should have the highest sequence numbers) so that the security appliance evaluates other crypto maps first. It examines the dynamic crypto map set only when the other (static) map entries do not match.

Similar to static crypto map sets, a dynamic crypto map set consists of all of the dynamic crypto maps with the same dynamic-map-name. The dynamic-seq-num differentiates the dynamic crypto maps in a set. If you configure a dynamic crypto map, insert a permit ACL to identify the data flow of the IPsec peer for the crypto access list. Otherwise the security appliance accepts any data flow identity the peer proposes.

**Caution**

Do not assign static (default) routes for traffic to be tunneled to a security appliance interface configured with a dynamic crypto map set. To identify the traffic that should be tunneled, add the ACLs to the dynamic crypto map. Use care to identify the proper address pools when configuring the ACLs associated with remote access tunnels. Use Reverse Route Injection to install routes only after the tunnel is up.

The procedure for using a dynamic crypto map entry is the same as the basic configuration described in “[Creating a Basic IPsec Configuration](#),” except that instead of creating a static crypto map, you create a dynamic crypto map entry. You can also combine static and dynamic map entries within a single crypto map set.

Create a crypto dynamic map entry as follows:

**Step 1** (Optional) Assign an access list to a dynamic crypto map:

```
crypto dynamic-map dynamic-map-name dynamic-seq-num match address access-list-name
```

This determines which traffic should be protected and not protected.

For example:

```
crypto dynamic-map dyn1 10 match address 101
```

In this example, access list 101 is assigned to dynamic crypto map “dyn1.” The map sequence number is 10.

**Step 2** Specify which transform sets are allowed for this dynamic crypto map. List multiple transform sets in order of priority (highest priority first).

```
crypto dynamic-map dynamic-map-name dynamic-seq-num set transform-set transform-set-name1,
[transform-set-name2, ...transform-set-name9]
```

For example:

```
crypto dynamic-map dyn 10 set transform-set myset1 myset2
```

In this example, when traffic matches access list 101, the SA can use either “myset1” (first priority) or “myset2” (second priority), depending on which transform set matches the transform sets of the peer.

- Step 3** (Optional) Specify the SA lifetime for the crypto dynamic map entry if you want to override the global lifetime value:

```
crypto dynamic-map dynamic-map-name dynamic-seq-num set security-association lifetime
{seconds seconds | kilobytes kilobytes}
```

For example:

```
crypto dynamic-map dyn1 10 set security-association lifetime seconds 2700
```

This example shortens the timed lifetime for dynamic crypto map “dyn1 10” to 2700 seconds (45 minutes). The time volume lifetime is not changed.

- Step 4** (Optional) Specify that IPsec ask for PFS when requesting new SAs for this dynamic crypto map, or should demand PFS in requests received from the peer:

```
crypto dynamic-map dynamic-map-name dynamic-seq-num set pfs [group1 | group2 | group5 |
group7]
```

For example:

```
crypto dynamic-map dyn1 10 set pfs group5
```

- Step 5** Add the dynamic crypto map set into a static crypto map set.

Be sure to set the crypto maps referencing dynamic maps to be the lowest priority entries (highest sequence numbers) in a crypto map set.

```
crypto map map-name seq-num ipsec-isakmp dynamic dynamic-map-name
```

For example:

```
crypto map mymap 200 ipsec-isakmp dynamic dyn1
```

---

## Providing Site-to-Site Redundancy

You can define multiple peers by using crypto maps to provide redundancy. This configuration is useful for site-to-site VPNs.

If one peer fails, the security appliance establishes a tunnel to the next peer associated with the crypto map. It sends data to the peer that it has successfully negotiated with, and that peer becomes the “active” peer. The “active” peer is the peer that the security appliance keeps trying first for follow-on negotiations until a negotiation fails. At that point the security appliance goes on to the next peer. The security appliance cycles back to the first peer when all peers associated with the crypto map have failed.

## Viewing an IPsec Configuration

[Table 27-5](#) lists commands you can enter to view information about your IPsec configuration.

**Table 27-5** *Commands to View IPsec Configuration Information*

| Command                                       | Purpose                                                                                                  |
|-----------------------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>show running-configuration crypto</b>      | Displays the entire crypto configuration, including IPsec, crypto maps, dynamic crypto maps, and ISAKMP. |
| <b>show running-config crypto ipsec</b>       | Displays the complete IPsec configuration.                                                               |
| <b>show running-config crypto isakmp</b>      | Displays the complete ISAKMP configuration.                                                              |
| <b>show running-config crypto map</b>         | Displays the complete crypto map configuration.                                                          |
| <b>show running-config crypto dynamic-map</b> | Displays the dynamic crypto map configuration.                                                           |
| <b>show all crypto map</b>                    | View all of the configuration parameters, including those with default values.                           |

## Clearing Security Associations

Certain configuration changes take effect only during the negotiation of subsequent SAs. If you want the new settings to take effect immediately, clear the existing SAs to reestablish them with the changed configuration. If the security appliance is actively processing IPsec traffic, clear only the portion of the SA database that the configuration changes affect. Reserve clearing the full SA database for large-scale changes, or when the security appliance is processing a small amount of IPsec traffic.

[Table 27-6](#) lists commands you can enter to clear and reinitialize IPsec SAs.

**Table 27-6** *Commands to Clear and Reinitialize IPsec SAs*

| Command                                     | Purpose                                                                                                |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------|
| <b>clear configure crypto</b>               | Removes an entire crypto configuration, including IPsec, crypto maps, dynamic crypto maps, and ISAKMP. |
| <b>clear configure crypto ca trustpoint</b> | Removes all trustpoints.                                                                               |
| <b>clear configure crypto dynamic-map</b>   | Removes all dynamic crypto maps. Includes keywords that let you remove specific dynamic crypto maps.   |
| <b>clear configure crypto map</b>           | Removes all crypto maps. Includes keywords that let you remove specific crypto maps.                   |
| <b>clear configure crypto isakmp</b>        | Removes the entire ISAKMP configuration.                                                               |
| <b>clear configure crypto isakmp policy</b> | Removes all ISAKMP policies or a specific policy.                                                      |
| <b>clear crypto isakmp sa</b>               | Removes the entire ISAKMP SA database.                                                                 |

## Clearing Crypto Map Configurations

The **clear configure crypto** command includes arguments that let you remove elements of the crypto configuration, including IPsec, crypto maps, dynamic crypto maps, CA trustpoints, all certificates, certificate map configurations, and ISAKMP.

Be aware that if you enter the **clear configure crypto** command without arguments, you remove the entire crypto configuration, including all certificates.


For more information, see the **clear configure crypto** command in the *Cisco Security Appliance Command Reference*.

## Supporting the Nokia VPN Client

The security appliance supports connections from Nokia VPN Clients on Nokia 92xx Communicator series phones using the Challenge/Response for Authenticated Cryptographic Keys (CRACK) protocol. CRACK is ideal for mobile IPsec-enabled clients that use legacy authentication techniques instead of digital certificates. It provides mutual authentication when the client uses a legacy based secret-key authentication technique such as RADIUS and the gateway uses public-key authentication.

The Nokia back-end services must be in place to support both Nokia clients and the CRACK protocol. This requirement includes the Nokia Security Services Manager (NSSM) and Nokia databases as shown in Figure 27-5.

**Figure 27-5** Nokia 92xx Communicator Service Requirement



To support the Nokia VPN Client, perform the following step on the security appliance:

- Enable CRACK authentication using the **crypto isakmp policy priority authentication** command with the **crack** keyword in global configuration mode. For example:

```
hostname(config)# crypto isakmp policy 2
```

```
hostname(config-isakmp-policy) # authentication crack
```

If you are using digital certificates for client authentication, perform the following additional steps:

- Step 1** Configure the trustpoint and remove the requirement for a fully qualified domain name. The trustpoint might be NSSM or some other CA. In this example, the trustpoint is named CompanyVPNCA:

```
hostname(config) # crypto ca trustpoint CompanyVPNCA
hostname(config-ca-trustpoint) # fqdn none
```

- Step 2** To configure the identity of the ISAKMP peer, perform one of the following steps:

- a. Use the **crypto isakmp identity** command with the **hostname** keyword. For example:

```
hostname(config) # crypto isakmp identity hostname
```

—or—

- b. Use the **crypto isakmp identity** command with the **auto** keyword to configure the identity to be automatically determined from the connection type. For example:

```
hostname(config) # crypto isakmp identity auto
```



**Note** If you use the **crypto isakmp identity auto** command, you must be sure that the DN attribute order in the client certificate is CN, OU, O, C, St, L.

To learn more about the Nokia services required to support the CRACK protocol on Nokia clients, and to ensure they are installed and configured properly, contact your local Nokia representative.







# CHAPTER 28

## Configuring L2TP over IPSec

---

This chapter describes how to configure IPSec over L2TP on the security appliance, and includes the following topics:

- [L2TP Overview, page 28-1](#)
- [Configuring L2TP over IPSec Connections, page 28-2](#)
- [Viewing L2TP over IPSec Connection Information, page 28-5](#)

### L2TP Overview

Layer 2 Tunneling Protocol (L2TP) is a VPN tunneling protocol which allows remote clients to use the public IP network to securely communicate with private corporate network servers. L2TP uses PPP over UDP (port 1701) to tunnel the data.

L2TP protocol is based on the client/server model. The function is divided between the L2TP Network Server (LNS), and the L2TP Access Concentrator (LAC). The LNS typically runs on a network gateway such as a router, while the LAC can be a dial-up Network Access Server (NAS), or a PC with a bundled L2TP client such as Microsoft Windows 2000.

The primary benefit of configuring L2TP with IPSec in a remote access scenario is that remote users can access a VPN over a public IP network without a gateway or a dedicated line, enabling remote access from virtually anywhere with POTS. An additional benefit is that the only client requirement for VPN access is the use of Windows 2000 with Microsoft Dial-Up Networking (DUN). No additional client software, such as Cisco VPN client software, is required.

To configure L2TP over IPSec, first configure IPSec transport mode to enable IPSec with L2TP. Then configure L2TP with a virtual private dial-up network VPDN group.

The configuration of L2TP with IPSec supports certificates using the pre-shared keys or RSA signature methods, and the use of dynamic (as opposed to static) crypto maps. This summary of tasks assumes completion of IKE, as well as pre-shared keys or RSA signature configuration. See “[Chapter 39, “Certificate Configuration,”](#)” for the steps to configure pre-shared keys, RSA, and dynamic crypto maps.



#### Note

L2TP with IPSec on the security appliance allows the LNS to interoperate with the Windows 2000 L2TP client. Interoperability with LACs from Cisco and other vendors is currently not supported. Only L2TP with IPSec is supported, native L2TP itself is not supported on security appliance.

The minimum IPSec security association lifetime supported by the Windows 2000 client is 300 seconds. If the lifetime on the security appliance is set to less than 300 seconds, the Windows 2000 client ignores it and replaces it with a 300 second lifetime.

## IPSec Transport and Tunnel Modes


By default, the security appliance uses IPSec tunnel mode—the entire original IP datagram is encrypted, and it becomes the payload in a new IP packet. This mode allows a network device, such as a router, to act as an IPSec proxy. That is, the router performs encryption on behalf of the hosts. The source router encrypts packets and forwards them along the IPSec tunnel. The destination router decrypts the original IP datagram and forwards it on to the destination system. The major advantage of tunnel mode is that the end systems do not need to be modified to receive the benefits of IPSec. Tunnel mode also protects against traffic analysis; with tunnel mode, an attacker can only determine the tunnel endpoints and not the true source and destination of the tunneled packets, even if they are the same as the tunnel endpoints.

However, the Windows 2000 L2TP/IPSec client uses IPSec transport mode—only the IP payload is encrypted, and the original IP headers are left intact. This mode has the advantages of adding only a few bytes to each packet and allowing devices on the public network to see the final source and destination of the packet. [Figure 28-1](#) illustrates the differences between IPSec Tunnel and Transport modes.

Therefore, In order for Windows 2000 L2TP/IPSec clients to connect to the security appliance, you must configure IPSec transport mode for a transform set using the **crypto ipsec transform-set trans\_name mode transport** command. This command is the configuration procedure that follows, “[Configuring L2TP over IPSec Connections](#)” section on page 28-2.

With this capability (transport), you can enable special processing (for example, QoS) on the intermediate network based on the information in the IP header. However, the Layer 4 header will be encrypted, limiting the examination of the packet. Unfortunately, transmitting the IP header in clear text, transport mode allows an attacker to perform some traffic analysis.

**Figure 28-1** *IPSec in Tunnel and Transport Modes*



## Configuring L2TP over IPSec Connections

To configure the security appliance to accept L2TP over IPSec connections, follow these steps:

**Note**

The security appliance does not establish an L2TP/IPSec tunnel with Windows 2000 if either the Cisco VPN Client Version 3.x or the Cisco VPN 3000 Client Version 2.5 is installed. Disable the *Cisco VPN Service* for the Cisco VPN Client Version 3.x, or the *ANetIKE Service* for the Cisco VPN 3000 Client Version 2.5 from the Services panel in Windows 2000 (click **Start>Programs>Administrative Tools>Services**). Then restart the IPSec Policy Agent Service from the **Services** panel, and reboot the machine.

- 
- Step 1** Specify IPSec to use transport mode rather than tunnel mode with the **mode** keyword of the **crypto ipsec transform-set** command:
- ```
hostname(config)# crypto ipsec transform-set trans_name mode transport
```
- Step 2** (Optional) Specify the local address pool used to allocate the IP address to the client using the **address-pool** command in tunnel-group general-attributes mode:
- ```
hostname(config)# tunnel-group name general-attributes
hostname(config-tunnel-general)# address-pool pool_name
```
- Step 3** (Optional) Instruct the security appliance to send DNS server IP addresses to the client with the **dns value** command from group policy configuration mode:
- ```
hostname(config)# group-policy group_policy_name attributes
hostname(config-group-policy)# dns value [none | IP_primary [IP_secondary]]
```
- Step 4** (Optional) Instruct the security appliance to send WINS server IP addresses to the client using the **wins-server** command from group policy configuration mode:
- ```
hostname(config-group-policy)# wins-server value [none | IP_primary [IP_secondary]]
```
- Step 5** (Optional) Generate a AAA accounting start and stop record for an L2TP session using the **accounting-server-group** command from tunnel group general-attributes mode:
- ```
hostname(config)# tunnel-group name general-attributes
hostname(config-tunnel-general)# accounting-server-group aaa_server_group
```
- Step 6** Configure L2TP over IPSec as a valid VPN tunneling protocol for a group or user with the **vpn-tunnel-protocol l2tp-ipsec** command:
- For a group, enter group-policy attributes mode:
- ```
hostname(config)# group-policy group_policy_name attributes
hostname(config-group-policy)# vpn-tunnel-protocol l2tp-ipsec
```
- For a user, enter username attributes mode:
- ```
hostname(config)# username user_name attributes
hostname(config-username)# vpn-tunnel-protocol l2tp-ipsec
```
- Step 7** Create a tunnel group with the **tunnel-group** command, and link the name of the group policy to the tunnel group with the **default-group-policy** command from tunnel group general-attributes mode:
- ```
hostname(config)# tunnel-group name type ipsec-ra
hostname(config)# tunnel-group name general-attributes
hostname(config-tunnel-general)# group-policy group_policy_name
```
- Step 8** Configure the PPP authentication protocol using the **authentication type** command from tunnel group ppp-attributes mode. [Table 28-1](#) shows the types of PPP authentication, and their characteristics.
- ```
hostname(config)# tunnel-group name ppp-attributes
hostname(config-ppp)# authentication pap
```

Table 28-1 Authentication Type Characteristics

Keyword	Authentication Type	Characteristics
chap	CHAP	In response to the server challenge, the client returns the encrypted [challenge plus password] with a cleartext username. This protocol is more secure than the PAP, but it does not encrypt data.
eap-proxy	EAP	Enables EAP which permits the security appliance to proxy the PPP authentication process to an external RADIUS authentication server.
ms-chap-v1 ms-chap-v2	Microsoft CHAP, Version 1 Microsoft CHAP, Version, 2	Similar to CHAP but more secure in that the server stores and compares only encrypted passwords rather than cleartext passwords as in CHAP. This protocol also generates a key for data encryption by MPPE.
pap	PAP	Passes cleartext username and password during authentication and is not secure.

- Step 9** Specify a method to authenticate users attempting L2TP over IPSec connections. Use the **authentication-server-group** command from tunnel-group general-attributes mode to configure the security appliance to use an authentication server or its own local database.

Using an Authentication Server

To use an authentication server, use the **authentication server group** keyword:

```
hostname(config)# tunnel-group name general-attributes
hostname(config-tunnel-general)# authentication-server-group auth_server_group
```

Using the Local Database

To use the local database, enter the **LOCAL** keyword.

```
hostname(config)# tunnel-group name general-attributes
hostname(config-tunnel-general)# authentication-server-group LOCAL
```



Note

The security appliance only supports the PPP authentications PAP and Microsoft CHAP, Versions 1 and 2, on the local database. EAP and CHAP are performed by proxy authentication servers. Therefore, if a remote user belongs to a tunnel group configured with the **authentication eap-proxy** or **authentication chap** commands, and the security appliance is configured to use the local database, that user will not be able to connect.

- Step 10** Create a user in the local database with the **username** command from global configuration mode.
- If the user is an L2TP client using Microsoft CHAP, Version 1 or Version 2, and the security appliance is configured to authenticate against the local database, you must include the **mschap** keyword. For Example:

```
hostname(config)# username t_wmith password eu5d93h mschap
```

- Step 11** Configure the interval (in seconds) between hello messages using the **l2tp tunnel hello** command in global configuration mode:

```
hostname(config)# l2tp tunnel hello seconds
```

Step 12 (Optional) If you expect multiple L2TP clients behind a NAT device to attempt L2TP over IPSec connections to the security appliance, you must enable NAT traversal so that ESP packets can pass through one or more NAT devices.

To enable NAT traversal globally, check that ISAKMP is enabled (you can enable it with the **crypto isakmp enable** command) in global configuration mode and then use the **crypto isakmp nat-traversal** command. For example:

```
hostname(config)# crypto isakmp enable
hostname(config)# crypto isakmp nat-traversal 30
```

Tunnel Group Switching

Tunnel Group Switching enables the security appliance to associate different users that are establishing L2TP over IPSec connections with different tunnel groups. Since each tunnel group has its own AAA server group and IP address pools, users can be authenticated through methods specific to their tunnel group.

With this feature, instead of sending just a username, the user sends a username and a group name in the format *username@group_name*, where “@” represents a delimiter that you can configure, and the group name is the name of a tunnel group that has been configured on the security appliance.

To enable Tunnel Group Switching, you must enable Strip Group processing using the **strip-group** command from tunnel-group general-attributes mode. When enabled, the security appliance selects the tunnel group for user connections by obtaining the group name from the username presented by the VPN client. The security appliance then sends only the user part of the username for authorization and authentication. Otherwise (if disabled), the security appliance sends the entire username, including the realm. In the following example, Strip Group processing is enabled for the tunnel-group *telecommuters*:

```
asa1(config)# tunnel-group telecommuters general-attributes
asa1(config-tunnel-general)# strip-group
```

Viewing L2TP over IPSec Connection Information

The **show vpn-sessiondb** command includes protocol filters that you can use to view detailed information about L2TP over IPSec connections. The full command from global configuration mode is **show vpn-sessiondb detailed remote filter protocol l2tpOverIpsec**.

The following example shows the details of a single L2TP over IPSec connection:

```
hostname# show vpn-sessiondb detail remote filter protocol l2tpOverIpsec
```

Session Type: Remote Detailed

```
Username      : b_smith
Index         : 1
Assigned IP   : 90.208.1.200      Public IP     : 70.208.1.212
Protocol      : L2TPOverIPSec    Encryption    : 3DES
Hashing       : SHA1
Bytes Tx      : 418464            Bytes Rx      : 424440
Client Type   :                  Client Ver    :
Group Policy  : DfltGrpPolicy
Tunnel Group  : DefaultRAGroup
Login Time    : 13:24:48 UTC Thu Mar 30 2006
Duration      : 1h:09m:18s
Filter Name   : #ACSACL#-IP-ACL4Clients-440fa5aa
NAC Result    : N/A
Posture Token :
```

```

IKE Sessions: 1
IPSec Sessions: 1
L2TPOverIPSec Sessions: 1

```

IKE:

```

Session ID      : 1
UDP Src Port    : 500
IKE Neg Mode    : Main
Encryption      : 3DES
Rekey Int (T)   : 28800 Seconds
D/H Group       : 2
UDP Dst Port    : 500
Auth Mode       : preSharedKeys
Hashing         : SHA1
Rekey Left(T)   : 24643 Seconds

```

IPSec:

```

Session ID      : 2
Local Addr      : 80.208.1.2/255.255.255.255/17/1701
Remote Addr     : 70.208.1.212/255.255.255.255/17/1701
Encryption      : 3DES
Encapsulation   : Transport
Hashing         : SHA1
Rekey Int (T)   : 3600 Seconds
Rekey Int (D)   : 95000 K-Bytes
Idle Time Out   : 30 Minutes
Bytes Tx        : 419064
Pkts Tx         : 4201
Rekey Left(T)   : 2856 Seconds
Rekey Left(D)   : 95000 K-Bytes
Idle TO Left    : 30 Minutes
Bytes Rx        : 425040
Pkts Rx         : 4227

```

L2TPOverIPSec:

```

Session ID      : 3
Username        : l2tp
Assigned IP     : 90.208.1.200
Encryption      : none
Idle Time Out   : 30 Minutes
Bytes Tx        : 301386
Pkts Tx         : 4198
Auth Mode       : PAP
Idle TO Left    : 30 Minutes
Bytes Rx        : 306480
Pkts Rx         : 4224

```

The following example shows the details of a single L2TP over IPSec over NAT connection:

```
hostname# show vpn-sessiondb detail remote filter protocol L2TPOverIPSecOverNatT
```

```
Session Type: Remote Detailed
```

```

Username        : v_gonzalez
Index           : 2
Assigned IP     : 90.208.1.202
Protocol        : L2TPOverIPSecOverNatT
Hashing         : MD5
Bytes Tx        : 1009
Client Type     :
Group Policy    : DfltGrpPolicy
Tunnel Group    : l2tpcert
Login Time      : 14:35:15 UTC Thu Mar 30 2006
Duration        : 0h:00m:07s
Filter Name     :
NAC Result      : N/A
Posture Token    :
Public IP       : 70.208.1.2
Encryption      : 3DES
Bytes Rx        : 2241
Client Ver      :

```

```

IKE Sessions: 1
IPSecOverNatT Sessions: 1
L2TPOverIPSecOverNatT Sessions: 1

```

```
IKE:
```

```

Session ID      : 1
UDP Src Port    : 4500
IKE Neg Mode    : Main
Encryption      : 3DES
Rekey Int (T)   : 300 Seconds
D/H Group      : 2
UDP Dst Port    : 4500
Auth Mode       : rsaCertificate
Hashing         : MD5
Rekey Left(T)   : 294 Seconds

IPSecOverNatT:
Session ID      : 2
Local Addr      : 80.208.1.2/255.255.255.255/17/1701
Remote Addr     : 70.208.1.2/255.255.255.255/17/0
Encryption      : 3DES
Hashing         : MD5
Encapsulation   : Transport
Rekey Int (T)   : 300 Seconds
Idle Time Out   : 1 Minutes
Bytes Tx        : 1209
Pkts Tx         : 20
Rekey Left(T)   : 293 Seconds
Idle TO Left    : 1 Minutes
Bytes Rx        : 2793
Pkts Rx         : 32

L2TPOverIPSecOverNatT:
Session ID      : 3
Username        : v_gonzalez
Assigned IP     : 90.208.1.202
Encryption      : none
Auth Mode       : PAP
Idle Time Out   : 1 Minutes
Idle TO Left    : 1 Minutes
Bytes Tx        : 584
Pkts Tx         : 18
Bytes Rx        : 2224
Pkts Rx         : 30
=====

```

Using L2TP Debug Commands

You can display L2TP debug information using the **debug l2tp** command in privileged EXEC mode. To disable the display of debug information, use the **no** form of this command:

debug l2tp {data | error | event | packet} level

data displays data packet trace information.

error displays error events.

event displays L2TP connection events.

packet displays packet trace information.

level sets the debug message level to display, between 1 and 255. The default is 1. To display additional messages at higher levels, set the level to a higher number.

The following example enables L2TP debug messages for connection events. The **show debug** command reveals that L2TP debug messages are enabled.

```

hostname# debug l2tp event 1
hostname# show debug
debug l2tp event enabled at level 1
hostname#

```

Enabling IPSec Debug

IPSec debug information can be added to a Windows 2000 client by adding the following registry:

Step 1 Run the Windows 2000 registry editor: REGEDIT.

- Step 2** Locate the following registry entry:
MyComputer\HKEY_LOCAL_MACHINE\CurrentControlSet\Services\PolicyAgent
- Step 3** Create the key by entering **oakley**.
- Step 4** Create the DWORD by entering **EnableLogging**.
- Step 5** Set the “Enable Logging” value to “1”.
- Step 6** Stop and Start the IPSec Policy Agent (click **Start>Programs>Administrative Tools>Services**). The debug file will be found at “%windir%\debug\oakley.log”.
-

Getting Additional Information

Additional information on various topics can be found at www.microsoft.com:

<http://support.microsoft.com/support/kb/articles/Q240/2/62.ASP>

How to Configure an L2TP/IPSec Connection Using Pre-Shared Keys Authentication:

<http://support.microsoft.com/support/kb/articles/Q253/4/98.ASP>

How to Install a Certificate for Use with IP Security (IPSec):

http://www.microsoft.com/windows2000/en/server/help/default.asp?url=/WINDOWS2000/en/server/help/sag_VPN_us26.htm

How to use a Windows 2000 Machine Certificate for L2TP over IPSec VPN Connections:

<http://www.microsoft.com/windows2000/techinfo/planning/security/ipsecsteps.asp#heading3>

How to Create a Custom MMC Console and Enabling Audit Policy for Your Computer:

<http://support.microsoft.com/support/kb/articles/Q259/3/35.ASP>



CHAPTER 29

Setting General IPSec VPN Parameters

The security appliance implementation of virtual private networking includes useful features that do not fit neatly into categories. This chapter describes some of these features. It includes the following sections:

- [Configuring VPNs in Single, Routed Mode, page 29-1](#)
- [Configuring IPSec to Bypass ACLs, page 29-1](#)
- [Permitting Intra-Interface Traffic, page 29-2](#)
- [Setting Maximum Active IPSec VPN Sessions, page 29-3](#)
- [Using Client Update to Ensure Acceptable Client Revision Levels, page 29-3](#)
- [Understanding Load Balancing, page 29-5](#)
- [Configuring Load Balancing, page 29-9](#)
- [Configuring VPN Session Limits, page 29-11](#)

Configuring VPNs in Single, Routed Mode

VPNs work only in single, routed mode. VPN functionality is unavailable in configurations that include either security contexts, also referred to as multi-mode firewall, or Active/Active stateful failover.

The exception to this caveat is that you can configure and use one connection for administrative purposes to (not through) the security appliance in transparent mode.

Configuring IPSec to Bypass ACLs

To permit any packets that come from an IPSec tunnel without checking ACLs for the source and destination interfaces, enter the **sysopt connection permit-ipsec** command in global configuration mode.

You might want to bypass interface ACLs for IPSec traffic if you use a separate VPN concentrator behind the security appliance and want to maximize the security appliance performance. Typically, you create an ACL that permits IPSec packets using the **access-list** command and apply it to the source interface. Using an ACL is more secure because you can specify the exact traffic you want to allow through the security appliance.

The syntax is **sysopt connection permit-ipsec**. The command has no keywords or arguments.

The following example enables IPSec traffic through the security appliance without checking ACLs:

```
hostname(config)# sysopt connection permit-ipsec
```

Permitting Intra-Interface Traffic

The security appliance includes a feature that lets a VPN client send IPSec-protected traffic to another VPN user by allowing such traffic in and out of the same interface. Also called “hairpinning”, this feature can be thought of as VPN spokes (clients) connecting through a VPN hub (security appliance).

In another application, this feature can redirect incoming VPN traffic back out through the same interface as unencrypted traffic. This would be useful, for example, to a VPN client that does not have split tunneling but needs to both access a VPN and browse the Web.

Figure 29-1 shows VPN Client 1 sending secure IPSec traffic to VPN Client 2 while also sending unencrypted traffic to a public Web server.

Figure 29-1 *VPN Client Using Intra-Interface Feature for Hairpinning*

To configure this feature, use the **same-security-traffic** command in global configuration mode with its **intra-interface** argument.

The command syntax is **same-security-traffic permit {inter-interface | intra-interface}**.

The following example shows how to enable intra-interface traffic:

```
hostname(config)# same-security-traffic permit intra-interface  
hostname(config)#
```



Note

You use the **same-security-traffic** command, but with the **inter-interface** argument, to permit communication between interfaces that have the same security level. This feature is not specific to IPSec connections. For more information, see the “Configuring Interface Parameters” chapter of this guide.

To use hairpinning, you must apply the proper NAT rules to the security appliance interface, as discussed in the following section.

NAT Considerations for Intra-Interface Traffic

For the security appliance to send unencrypted traffic back out through the interface, you must enable NAT for the interface so that publicly routable addresses replace your private IP addresses (unless you already use public IP addresses in your local IP address pool). The following example applies an interface PAT rule to traffic sourced from the client IP pool:

```
hostname(config)# ip local pool clientpool 192.168.0.10-192.168.0.100
hostname(config)# global (outside) 1 interface
hostname(config)# nat (outside) 1 192.168.0.0 255.255.255.0
```

When the security appliance sends encrypted VPN traffic back out this same interface, however, NAT is optional. The VPN-to-VPN hairpinning works with or without NAT. To apply NAT to all outgoing traffic, implement only the commands above. To exempt the VPN-to-VPN traffic from NAT, add commands (to the example above) that implement NAT exemption for VPN-to-VPN traffic, such as:

```
hostname(config)# access-list nonat permit ip 192.168.0.0 255.255.255.0 192.168.0.0
255.255.255.0
hostname(config)# nat (outside) 0 access-list nonat
```

For more information on NAT rules, see the “Applying NAT” chapter of this guide.

Setting Maximum Active IPsec VPN Sessions

To limit VPN sessions to a lower value than the security appliance allows, enter the **vpn-sessiondb max-session-limit** command in global configuration mode.

- This command applies to all types of VPN sessions, including WebVPN.
- This limit affects the calculated load percentage for VPN Load Balancing.

The syntax is **vpn-sessiondb max-session-limit {session-limit}**.

The following example shows how to set a maximum VPN session limit of 450:

```
hostname (config)# vpn-sessiondb max-session-limit 450
hostname (config)#
```

Using Client Update to Ensure Acceptable Client Revision Levels

The client update feature lets administrators at a central location automatically notify VPN client users that it is time to update the VPN client software and the VPN 3002 hardware client image.

Remote users might be using outdated VPN software or hardware client versions. You can use the **client-update** command at any time to enable updating client revisions; specify the types and revision numbers of clients to which the update applies; provide a URL or IP address from which to get the update; and, in the case of Windows clients, optionally notify users that they should update their VPN client version. For Windows clients, you can provide a mechanism for users to accomplish that update. For VPN 3002 hardware client users, the update occurs automatically, with no notification. This command applies only to the IPsec remote-access tunnel-group type.

To perform client update, enter the **client-update** command in either general configuration mode or tunnel-group ipsec-attributes configuration mode. If the client is already running a software version on the list of revision numbers, it does not need to update its software. If the client is not running a software version on the list, it should update. The following procedure tells how to perform a client-update:

Step 1 In global configuration mode, enable client update by entering the command:

```
hostname(config)# client-update enable
hostname(config)#
```

Step 2 In global configuration mode, specify the parameters for the client update that you want to apply to all clients of a particular type. That is, specify the type of client, the URL or IP address from which to get the updated image, and the acceptable revision number or numbers for that client. You can specify up to four revision numbers, separated by commas.

If the user's client revision number matches one of the specified revision numbers, there is no need to update the client. This command specifies the client-update values for all clients of the specified type across the entire security appliance.

The syntax of the command to do this is:

```
hostname(config)# client-update type type url url-string rev-nums rev-numbers
hostname(config)#
```

The available client types are **win9X** (includes Windows 95, Windows 98 and Windows ME platforms), **winnt** (includes Windows NT 4.0, Windows 2000 and Windows XP platforms), **windows** (Includes all Windows based platforms), and **vpn3002** (VPN 3002 hardware client).

If the client is already running a software version on the list of revision numbers, it does not need to update its software. If the client is not running a software version on the list, it should update. You can specify up to three of these client update entries. The keyword **windows** covers all of the allowable Windows platforms. If you specify **windows**, do not specify the individual Windows client types.



Note

For all Windows clients, you must use the protocol `http://` or `https://` as the prefix for the URL. For the VPN 3002 hardware client, you must specify protocol `tftp://` instead.

The following example configures client update parameters for the remote-access tunnel-group. It designates the revision number, 4.6.1 and the URL for retrieving the update, which is `https://support/updates:`

```
hostname(config)# client-update type windows url https://support/updates/ rev-nums 4.6.1
hostname(config)#
```

Alternatively, you can configure client update just for individual tunnel-groups, rather than for all clients of a particular type. (See Step 3.)

VPN 3002 clients update without user intervention and users receive no notification message. The following example applies only to VPN 3002 Hardware Clients. Entered in tunnel-group ipsec-attributes configuration mode, it configures client update parameters for the IPSec remote-access tunnel-group "salesgrp". It designates the revision number, 4.7 and uses the TFTP protocol for retrieving the updated software from the site with the IP address 192.168.1.1:

```
hostname(config)# tunnel-group salesgrp type ipsec-ra
hostname(config)# tunnel-group salesgrp ipsec-attributes
hostname(config-tunnel-ipsec)# client-update type vpn3002 url tftp:192.168.1.1 rev-nums 4.7
hostname(config-tunnel-ipsec)#
```

**Note**

You can have the browser automatically start an application by including the application name at the end of the URL; for example: **https://support/updates/vpnclient.exe**.

Step 3

To define a set of client-update parameters for a particular ipsec-ra tunnel group, do the following. In tunnel-group ipsec-attributes mode, specify the tunnel-group name and its type, the URL or IP address from which to get the updated image, and a revision number. If the user's client's revision number matches one of the specified revision numbers, there is no need to update the client; for example, for a Windows client:

```
hostname(config)# tunnel-group remotegrp type ipsec-ra
hostname(config)# tunnel-group remotegrp ipsec-attributes
hostname(config-tunnel-ipsec)# client-update type windows url https://support/updates/
rev-nums 4.6.1
hostname(config-tunnel-ipsec)#
```

Step 4

Optionally, you can send a notice to active users with outdated Windows clients that their client needs updating. For these users, a pop-up window appears, offering them the opportunity to launch a browser and download the updated software from the site that you specified in the URL. The only part of this message that you can configure is the URL. (See Step 2 or 3.) Users who are not active get a notification message the next time they log on. You can send this notice to all active clients on all tunnel groups, or you can send it to clients on a particular tunnel group. For example, to notify all active clients on all tunnel groups, you would enter the following command in privileged EXEC mode:

```
hostname# client-update all
hostname#
```

If the user's client's revision number matches one of the specified revision numbers, there is no need to update the client, and no notification message is sent to the user. VPN 3002 clients update without user intervention and users receive no notification message.

**Note**

If you specify the client-update type as **windows** (specifying all Windows-based platforms) and later want to enter a client-update type of **win9x** or **winnt** for the same entity, you must first remove the windows client type with the **no** form of the command, then use new client-update commands to specify the new client types.

Understanding Load Balancing

If you have a remote-access configuration in which you are using two or more security appliances or VPN Concentrators connected on the same network to handle remote sessions, you can configure these devices to share their session load. This feature is called *load balancing*. To implement load balancing, you group together logically two or more devices on the same private LAN-to-LAN network, private subnet, and public subnet into a *virtual cluster*.

All devices in the virtual cluster carry session loads. Load balancing directs session traffic to the least loaded device in the cluster, thus distributing the load among all devices. It makes efficient use of system resources and provides increased performance and high availability.

One device in the virtual cluster, the *virtual cluster master*, directs incoming traffic to the other devices, called *secondary devices*. The virtual cluster master monitors all devices in the cluster, keeps track of how busy each is, and distributes the session load accordingly. The role of virtual cluster master is not

tied to a physical device; it can shift among devices. For example, if the current virtual cluster master fails, one of the secondary devices in the cluster takes over that role and immediately becomes the new virtual cluster master.

**Note**

The output of a **show** command might show the secondary devices in the cluster as backup devices.

The virtual cluster appears to outside clients as a single *virtual cluster IP address*. This IP address is not tied to a specific physical device. It belongs to the current virtual cluster master; hence, it is virtual. A VPN Client attempting to establish a connection connects first to this virtual cluster IP address. The virtual cluster master then sends back to the client the public IP address of the least-loaded available host in the cluster. In a second transaction (transparent to the user), the client connects directly to that host. In this way, the virtual cluster master directs traffic evenly and efficiently across resources.

**Note**

All clients other than the Cisco VPN Client or the Cisco 3002 Hardware Client should connect directly to the security appliance as usual; they do not use the virtual cluster IP address.

If a machine in the cluster fails, the terminated sessions can immediately reconnect to the virtual cluster IP address. The virtual cluster master then directs these connections to another active device in the cluster. Should the virtual cluster master itself fail, a secondary device in the cluster immediately and automatically takes over as the new virtual session master. Even if several devices in the cluster fail, users can continue to connect to the cluster as long as any one device in the cluster is up and available.

Implementing Load Balancing

Enabling load balancing involves:

- Configuring the load-balancing cluster by establishing a common virtual cluster IP address, UDP port (if necessary), and IPSec shared secret for the cluster. These values should be configured identically for every device in the cluster.
- Configuring the participating device by enabling load balancing on the device and defining device-specific properties. These values vary from device to device.

**Note**

VPN load balancing requires an active 3DES/AES license. The security appliance checks for the existence of this crypto license before enabling load balancing. If it does not detect an active 3DES or AES license, the security appliance prevents the enabling of load balancing and also prevents internal configuration of 3DES by the load balancing system unless the license permits this usage.

Prerequisites

Load balancing is disabled by default. You must explicitly enable load balancing.

You must have first configured the public (outside) and private (inside) interfaces and also have previously configured the interface to which the virtual cluster IP address refers. You can use the **interface** and **nameif** commands to configure different names for these interfaces. Subsequent references in this section use the names outside and inside.

All devices that participate in a cluster must share the same cluster-specific values: IP address, encryption settings, encryption key, and port.

Eligible Platforms

A load-balancing cluster can include security appliance models ASA 5520 and above. You can also include VPN 3000 Series Concentrators in the cluster. While mixed configurations are possible, administration is generally simpler if the cluster is homogeneous.

Eligible Clients

Load balancing is effective only on remote sessions initiated with the following clients:

- Cisco VPN Client (Release 3.0 and later)
- Cisco VPN 3002 Hardware Client (Release 3.5 or later)
- Cisco PIX 501/506E when acting as an Easy VPN client.

Load balancing works with both IPSec clients and WebVPN sessions. All other clients, including LAN-to-LAN connections, can connect to a security appliance on which load balancing is enabled, but they cannot participate in load balancing.

VPN Load-Balancing Cluster Configurations

A load-balancing cluster can consist of all ASA Release 7.0(x) security appliances, all ASA Release 7.1(1) security appliances, all VPN 3000 Concentrators, or a mixture of these, subject to the following restrictions:

- Load-balancing clusters that consist of all ASA 7.0(x) security appliances, all ASA 7.1(1) security appliances, or all VPN 3000 Concentrators can run load balancing for a mixture of IPSec and WebVPN sessions.
- Load-balancing clusters that consist of a both of ASA 7.0(x) security appliances and VPN 3000 Concentrators can run load balancing for a mixture of IPSec and WebVPN sessions.
- Load-balancing clusters that include ASA 7.1(1) security appliances and either ASA 7.0(x) or VPN 3000 Concentrators or both can support only IPSec sessions. In such a configuration, however, the ASA 7.1(1) security appliances might not reach their full IPSec capacity. [“Scenario 1: Mixed Cluster with No WebVPN Connections” on page 8](#), illustrates this situation.

With Release 7.1(1), IPSec and WebVPN sessions count or weigh equally in determining the load that each device in the cluster carries. This represents a departure from the load balancing calculation for the ASA Release 7.0(x) software and the VPN 3000 Concentrator, in that these platforms both use a weighting algorithm that, on some hardware platforms, calculates WebVPN session load differently from IPSec session load.

The virtual master of the cluster assigns session requests to the members of the cluster. An ASA Release 7.1(1) security appliance regards all sessions, WebVPN or IPSec, as equal and assigns them accordingly. An ASA Release 7.0(x) security appliance or a VPN 3000 Concentrator performs a weighting calculation in assigning session loads.



Note

You can configure the number of IPSec and WebVPN sessions to allow, up to the maximum allowed by your configuration and license. See [Configuring VPN Session Limits, page 29-11](#) for a description of how to set these limits.

Some Typical Mixed Cluster Scenarios

If you have a mixed configuration—that is, if your load-balancing cluster includes devices running a mixture of ASA software releases or at least one security appliance running ASA Release 7.1(1) and a VPN 3000 Concentrator—the difference in weighting algorithms becomes an issue if the initial cluster master fails and another device takes over as master.

The following scenarios illustrate the use of VPN load balancing in clusters consisting of a mixture of security appliances running ASA Release 7.1(1) and ASA Release 7.0(x) software, as well as VPN 3000 Series Concentrators.

Scenario 1: Mixed Cluster with No WebVPN Connections

In this scenario, the cluster consists of a mixture of security appliances and VPN 3000 Concentrators. Some of the security appliance cluster peers are running ASA Release 7.0(x), and some are running Release 7.1(1). The pre-7.1(1) and VPN 3000 peers do not have any SSL VPN connections, and the 7.1(1) cluster peers have only the base SSL VPN license, which allows two WebVPN sessions, but there are no SSL VPN connections. In this case, all the connections are IPSec, and load balancing works fine.

The two WebVPN licenses have a very small effect on the user's taking advantage of the maximum IPSec session limit, and then only when a VPN 3000 Concentrator is the cluster master. In general, the smaller the number of WebVPN licenses is on a security appliance in a mixed cluster, the smaller the effect on the ASA 7.1(1) device being able to reach its IPSec session limit in a scenario where there are only IPSec sessions.

Scenario 2: Mixed Cluster Handling WebVPN Connections

Suppose, for example, a security appliance running ASA Release 7.1(1) software is the initial cluster master; then that device fails. Another device in the cluster takes over automatically as master and applies its own load-balancing algorithm to determine processor loads within the cluster. A cluster master running ASA Release 7.1(1) software cannot weight session loads in any way other than what that software provides. Therefore, it cannot assign a combination of IPSec and WebVPN session loads properly to ASA devices running earlier versions nor to VPN 3000 Concentrators. Conversely, a VPN 3000 Concentrator acting as the cluster master cannot assign loads properly to an ASA Release 7.1(1) security appliance. The following scenario illustrates this dilemma.

This scenario is similar to the previous one, in that the cluster consists of a mixture of security appliances and VPN 3000 Concentrators. Some of the security appliance cluster peers are running ASA Release 7.0(x) and some are running Release 7.1(1). In this case, however, the cluster is handling SSL VPN connections as well as IPSec connections.

If a device that is running software earlier than ASA Release 7.1(1) is the cluster master, the master applies the protocol and logic in effect prior to Release 7.1(1). That is, sessions might be directed to load-balancing peers that have exceeded their session limit. In that case, the user is denied access.

If the cluster master is a device running ASA Release 7.0(x) software, the old session-weighting algorithm applies only to the pre-7.1(1) peers in the cluster. No one should be denied access in this case. Because the pre-7.1(1) peers use the session-weighting algorithm, they are more lightly loaded.

An issue arises, however, because you cannot guarantee that the 7.1(1) peer is always the cluster master. If the cluster master fails, another peer assumes the role of master. The new master might be any of the eligible peers. Because of the innately unpredictability of the results, we recommend that you avoid configuring this type of cluster.

Configuring Load Balancing

To use load balancing, configure the following elements for each device that participates in the cluster.

- Public and private interfaces
- VPN load-balancing cluster attributes

**Note**

All participants in the cluster must have an identical cluster configuration, except for the device priority within the cluster.

Configuring the Public and Private Interfaces for Load Balancing

To configure the public (outside) and private (inside) interfaces for the load-balancing cluster devices, do the following steps:

- Step 1** Configure the public interface on the security appliance by entering the **interface** command with the **lbpublic** keyword in vpn-load-balancing configuration mode. This command specifies the name or IP address of the public interface for load balancing for this device:

```
hostname(config)# vpn load-balancing  
hostname(config-load-balancing)# interface lbpublic outside  
hostname(config-load-balancing)#
```

- Step 2** Configure the private interface on the security appliance by entering the **interface** command with the **lbprivate** keyword in vpn-load-balancing configuration mode. This command specifies the name or IP address of the private interface for load balancing for this device:

```
hostname(config-load-balancing)# interface lbprivate inside  
hostname(config-load-balancing)#
```

- Step 3** Set the priority to assign to this device within the cluster. The range is from 1 to 10. The priority indicates the likelihood of this device becoming the virtual cluster master, either at start-up or when an existing master fails. The higher you set the priority (for example, 10), the more likely it is that this device becomes the virtual cluster master.

```
hostname(config-load-balancing)# priority number  
hostname(config-load-balancing)#
```

For example, to assign this device a priority of 6 within the cluster, enter the following command:

```
hostname(config-load-balancing)# priority 6  
hostname(config-load-balancing)#
```

- Step 4** If you want to apply network address translation for this device, enter the **nat** command with the NAT assigned address for the device:

```
hostname(config-load-balancing)# nat ip_address  
hostname(config-load-balancing)#
```

For example, to assign this device a NAT address of 192.168.30.3, enter the following command:

```
hostname(config-load-balancing)# nat 192.168.30.3  
hostname(config-load-balancing)#
```

Configuring the Load Balancing Cluster Attributes

To configure the load-balancing cluster attributes for each device in the cluster, do the following steps:

- Step 1** Set up VPN load balancing by entering the `vpn load-balancing` command in global configuration mode:

```
hostname(config)# vpn load-balancing  
hostname(config-load-balancing)#
```

This enters `vpn-load-balancing` configuration mode, in which you can configure the remaining load-balancing attributes.

- Step 2** Configure the IP address of the cluster to which this device belongs. This command specifies the single IP address that represents the entire virtual cluster. Choose an IP address that is within the public subnet address range shared by all the security appliances in the virtual cluster

```
hostname(config-load-balancing)# cluster ip address ip_address  
hostname(config-load-balancing)#
```

For example, to set the cluster IP address to 192.168.10.10, enter the following command:

```
hostname(config-load-balancing)# cluster ip address 192.168.10.10  
hostname(config-load-balancing)#
```

- Step 3** Configure the cluster port. This command specifies the UDP port for the virtual cluster in which this device is participating. The default value is 9023. If another application is using this port, enter the UDP destination port number you want to use for load balancing.

```
hostname(config-load-balancing)# cluster port port_number  
hostname(config-load-balancing)#
```

For example, to set the cluster port to 4444, enter the following command:

```
hostname(config-load-balancing)# cluster port 4444  
hostname(config-load-balancing)#
```

- Step 4** Optionally, enable IPSec encryption for the cluster. The default is no encryption. This command enables or disables IPSec encryption. If you configure this check attribute, you must first specify and verify a shared secret. The security appliances in the virtual cluster communicate via LAN-to-LAN tunnels using IPSec. To ensure that all load-balancing information communicated between the devices is encrypted, enable this attribute.

```
hostname(config-load-balancing)# cluster encryption  
hostname(config-load-balancing)#
```



Note

When using encryption, you must have previously configured the load-balancing inside interface. If that interface is not enabled on the load-balancing inside interface, you get an error message when you try to configure cluster encryption.

If the load-balancing inside interface was enabled when you configured cluster encryption, but was disabled before you configured the participation of the device in the virtual cluster, you get an error message when you enter the **participate** command (or, in ASDM, select the Participate in Load Balancing Cluster check box), and encryption is not enabled for the cluster.

To use cluster encryption, you must enable `isakmp` on the inside interface, using the **crypto isakmp enable** command with the inside interface specified.

- Step 5** If you enable cluster encryption, you must also specify the IPsec shared secret by entering the **cluster key** command. This command specifies the shared secret to between IPsec peers when you have enabled IPsec encryption. The value you enter in the box appears as consecutive asterisk characters

```
hostname(config-load-balancing)# cluster key shared_secret
hostname(config-load-balancing)#
```

For example, to set the shared secret to 123456789, enter the following command:

```
hostname(config-load-balancing)# cluster key 123456789
hostname(config-load-balancing)#
```

- Step 6** Enable this device's participation in the cluster by entering the **participate** command:

```
hostname(config-load-balancing)# participate
hostname(config-load-balancing)#
```

Configuring VPN Session Limits

You can run as many IPsec and WebVPN sessions as your platform and license for the security appliance supports. To view the licensing information for your security appliance, enter the **show version** command in global configuration mode. The following example shows the command and the licensing information excerpted from the output of this command:

```
hostname(config)# show version
```

```
Cisco Adaptive Security Appliance Software Version 7.1(0)182
Device Manager Version 5.1(0)128
```

```
Licensed features for this platform:
Maximum Physical Interfaces : Unlimited
Maximum VLANs               : 100
Inside Hosts                 : Unlimited
Failover                     : Active/Active
VPN-DES                      : Enabled
VPN-3DES-AES                 : Enabled
Security Contexts            : 10
GTP/GPRS                     : Enabled
VPN Peers                    : 750
WebVPN Peers                  : 500
```

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

To limit the maximum number of active IPsec VPN sessions to a lower value than the security appliance allows, enter the **vpn-sessiondb max-session-limit** command in global configuration mode. This limit affects the calculated load percentage for VPN Load Balancing.

```
hostname(config)# vpn-sessiondb max-session-limit number_of_sessions
hostname(config)#
```

For example, if the security appliance license allows 750 IPsec sessions, and you want to limit the number of IPsec sessions to 500, enter the following command:

```
hostname(config)# vpn-sessiondb max-session-limit 500
hostname(config)#
```

To remove the session limit, use the **no** version of this command.:

```
hostname(config)# no vpn-sessiondb max-session-limit  
hostname(config)#
```

To limit WebVPN sessions to a lower value than the security appliance allows, use the **vpn-sessiondb max-webvpn-session-limit** command in global configuration mode. To remove the session limit, use the **no** version of this command.

```
hostname(config)# vpn-sessiondb max-webvpn-session-limit number_of_sessions  
hostname(config)#
```

For example, if the security appliance license allows 500 WebVPN sessions, and you want to limit the number of WebVPN sessions to 250, enter the following command:

```
hostname(config)# vpn-sessiondb max-webvpn-session-limit 250  
hostname(config)#
```

To remove the session limit, use the **no** version of this command.:

```
hostname(config)# no vpn-sessiondb max-webvpn-session-limit  
hostname(config)#
```

For a complete description of the features available with each license, see Appendix A, Feature Licenses and Specifications.



CHAPTER 30

Configuring Tunnel Groups, Group Policies, and Users

This chapter describes how to configure VPN tunnel groups, group policies, and users. This chapter includes the following sections.

- [Overview of Tunnel Groups, Group Policies, and Users, page 30-1](#)
- [Configuring Tunnel Groups, page 30-5](#)
- [Group Policies, page 30-31](#)
- [Configuring User Attributes, page 30-70](#)

In summary, you first configure tunnel groups to set the values for the connection. Then you configure group policies. These set values for users in the aggregate. Then you configure users, which can inherit values from groups and configure certain values on an individual user basis. This chapter describes how and why to configure these entities.

Overview of Tunnel Groups, Group Policies, and Users

Groups and users are core concepts in managing the security of virtual private networks (VPNs) and in configuring the security appliance. They specify attributes that determine user access to and use of the VPN. A *group* is a collection of users treated as a single entity. *Users* get their attributes from *group policies*. *Tunnel groups* identify the group policy for a specific connection. If you do not assign a particular group policy to a user, the default group policy for the connection applies.

Tunnel groups and group policies simplify system management. To streamline the configuration task, the security appliance provides a default LAN-to-LAN tunnel group, a default remote access tunnel group, a default WebVPN tunnel group, and a default group policy (DfltGrpPolicy). The default tunnel groups and group policy provide settings that are likely to be common for many users. As you add users, you can specify that they “inherit” parameters from a group policy. Thus you can quickly configure VPN access for large numbers of users.

If you decide to grant identical rights to all VPN users, then you do not need to configure specific tunnel groups or group policies, but VPNs seldom work that way. For example, you might allow a finance group to access one part of a private network, a customer support group to access another part, and an MIS group to access other parts. In addition, you might allow specific users within MIS to access systems that other MIS users cannot access. Tunnel groups and group policies provide the flexibility to do so securely.

**Note**

The security appliance also includes the concept of object groups, which are a superset of network lists. Object groups let you define VPN access to ports as well as networks. Object groups relate to ACLs rather than to group policies and tunnel groups. For more information about using object groups, see [Chapter 16, “Identifying Traffic with Access Lists.”](#)

Tunnel Groups

A tunnel group consists of a set of records that determines tunnel connection policies. These records identify the servers to which the tunnel user is authenticated, as well as the accounting servers, if any, to which connection information is sent. They also identify a default group policy for the connection, and they contain protocol-specific connection parameters. Tunnel groups include a small number of attributes that pertain to creating the tunnel itself. Tunnel groups include a pointer to a group policy that defines user-oriented attributes.

The security appliance provides the following default tunnel groups: DefaultL2Lgroup for LAN-to-LAN connections, DefaultRAGroup for remote access connections, and DefaultWEBVPNGroup for WebVPN connections. You can modify these default tunnel groups, but you cannot delete them. You can also create one or more tunnel groups specific to your environment. Tunnel groups are local to the security appliance and are not configurable on external servers.

Tunnel groups specify the following attributes:

- [General Tunnel-Group Connection Parameters, page 30-2](#)
- [IPSec Tunnel-Group Connection Parameters, page 30-3](#)
- [WebVPN Tunnel-Group Connection Parameters, page 30-4](#)

General Tunnel-Group Connection Parameters

General parameters are common to both IPSec and WebVPN connections. The general parameters include the following:

- Tunnel group name—You specify a tunnel-group name when you add or edit a tunnel group. The following considerations apply:
 - For clients that use preshared keys to authenticate, the tunnel group name is the same as the group name that an IPSec client passes to the security appliance.
 - Clients that use certificates to authenticate pass this name as part of the certificate, and the security appliance extracts the name from the certificate.
- Connection type—Connection types include IPSec remote access, IPSec LAN-to-LAN, and WebVPN. A tunnel group can have only one connection type.
- Authentication, Authorization, and Accounting servers—These parameters identify the server groups or lists that the security appliance uses for the following purposes:
 - Authenticating users
 - Obtaining information about services users are authorized to access
 - Storing accounting records

A server group can consist of one or more servers.

- Default group policy for the connection—A group policy is a set of user-oriented attributes. The default group policy is the group policy whose attributes the security appliance uses as defaults when authenticating or authorizing a tunnel user.
- Client address assignment method—This method includes values for one or more DHCP servers or address pools that the security appliance assigns to clients.
- Override account disabled—This parameter lets you override the “account-disabled” indicator received from a AAA server.
- Password management—This parameter lets you warn a user that the current password is due to expire in a specified number of days (the default is 14 days), then offer the user the opportunity to change the password.
- Strip group and strip realm—These parameters direct the way the security appliance processes the usernames it receives. They apply only to usernames received in the form user@realm. A realm is an administrative domain appended to a username with the @ delimiter (user@abc).

When you specify the **strip-group** command, the security appliance selects the tunnel group for user connections by obtaining the group name from the username presented by the VPN client. The security appliance then sends only the user part of the username for authorization/authentication. Otherwise (if disabled), the security appliance sends the entire username, including the realm.

Strip-realm processing removes the realm from the username when sending the username to the authentication or authorization server. If the command is enabled, the security appliance sends only the user part of the username authorization/authentication. Otherwise, the security appliance sends the entire username.

- Authorization required—This parameter lets you require authorization before a user can connect, or turn off that requirement.
- Authorization DN attributes—This parameter specifies which Distinguished Name attributes to use when performing authorization.

IPSec Tunnel-Group Connection Parameters

IPSec parameters include the following:

- A client authentication method: preshared keys, certificates, or both.
 - For IKE connections based on preshared keys, the alphanumeric key itself (up to 128 characters long), associated with the connection policy.
 - Peer-ID validation requirement—This parameter specifies whether to require validating the identity of the peer using the peer’s certificate.
- An extended hybrid authentication method: XAUTH and hybrid XAUTH.

You use **isakmp ikev1-user-authentication** command to implement hybrid XAUTH authentication when you need to use digital certificates for security appliance authentication and a different, legacy method for remote VPN user authentication, such as RADIUS, TACACS+ or SecurID.

- ISAKMP (IKE) keepalive settings. This feature lets the security appliance monitor the continued presence of a remote peer and report its own presence to that peer. If the peer becomes unresponsive, the security appliance removes the connection. Enabling IKE keepalives prevents hung connections when the IKE peer loses connectivity.

There are various forms of IKE keepalives. For this feature to work, both the security appliance and its remote peer must support a common form. This feature works with the following peers:

- Cisco VPN client (Release 3.0 and above)

- Cisco VPN 3000 Client (Release 2.x)
- Cisco VPN 3002 Hardware Client
- Cisco VPN 3000 Series Concentrators
- Cisco IOS software
- Cisco Secure PIX Firewall

Non-Cisco VPN clients do not support IKE keepalives.

If you are configuring a group of mixed peers, and some of those peers support IKE keepalives and others do not, enable IKE keepalives for the entire group. The feature does not affect the peers that do not support it.

If you disable IKE keepalives, connections with unresponsive peers remain active until they time out, so we recommend that you keep your idle timeout short. To change your idle timeout, see [“Configuring Group Policies” section on page 30-34](#).

**Note**

To reduce connectivity costs, disable IKE keepalives if this group includes any clients connecting via ISDN lines. ISDN connections normally disconnect if idle, but the IKE keepalive mechanism prevents connections from idling and therefore from disconnecting.

If you do disable IKE keepalives, the client disconnects only when either its IKE or IPSec keys expire. Failed traffic does not disconnect the tunnel with the Peer Timeout Profile values as it does when IKE keepalives are enabled.

**Note**

If you have a LAN-to-LAN configuration using IKE main mode, make sure that the two peers have the same IKE keepalive configuration. Both peers must have IKE keepalives enabled or both peers must have it disabled.

- If you configure authentication using digital certificates, you can specify whether to send the entire certificate chain (which sends the peer the identity certificate and all issuing certificates) or just the issuing certificates (including the root certificate and any subordinate CA certificates).
- You can notify users who are using outdated versions of Windows client software that they need to update their client, and you can provide a mechanism for them to get the updated client version. For VPN 3002 hardware client users, you can trigger an automatic update. You can configure and change the client-update, either for all tunnel groups or for particular tunnel groups.
- If you configure authentication using digital certificates, you can specify the name of the trustpoint that identifies the certificate to send to the IKE peer.

WebVPN Tunnel-Group Connection Parameters

The following attributes are specific to WebVPN connections:

- The authentication method, either AAA or certificate.
- The name of the customization to apply. Customizations determine the appearance of the windows that the user sees upon login. You configure the customization parameters as part of configuring WebVPN.
- The DNS server-group name. The DNS server group specifies the DNS server name, domain name, name server, number of retries, and timeout values for a DNS server to use for a tunnel group.

- One or more group aliases; these are alternate names by which the server can refer to a tunnel group. At login, the user selects the group name from a dropdown menu.
- One or more group URLs. If you configure this parameter, users coming in on a specified URL need not select a group at login.
- A group policy that grants a WebVPN user access rights that are different from the default group policy.
- The name of the NetBIOS Name Service server (nbns-server) to use for CIFS name resolution.

Configuring Tunnel Groups

The following sections describe the contents and configuration of tunnel groups:

- [Default IPsec Remote Access Tunnel Group Configuration, page 30-5](#)
- [Specifying a Name and Type for the IPsec Remote Access Tunnel Group, page 30-6](#)
- [Configuring IPsec Remote-Access Tunnel Groups, page 30-6](#)
- [Configuring LAN-to-LAN Tunnel Groups, page 30-13](#)
- [Configuring WebVPN Tunnel Groups, page 30-17](#)
- [Customizing Login Windows for WebVPN Users, page 30-23](#)

You can modify the default tunnel groups, and you can configure a new tunnel group as any of the three tunnel-group types. If you don't explicitly configure an attribute in a tunnel group, that attribute gets its value from the default tunnel group. The default tunnel-group type is ipsec-ra. The subsequent parameters depend upon your choice of tunnel type. To see the current configured and default configuration of all your tunnel groups, including the default tunnel group, enter the **show running-config all tunnel-group** command.

Maximum Tunnel Groups

The maximum number of tunnel groups that a security appliance can support is a function of the maximum number of concurrent VPN sessions for the platform + 5. For example, an ASA5505 can support a maximum of 25 concurrent VPN sessions allowing for 30 tunnel groups (25+5). Attempting to add an additional tunnel group beyond the limit results in the following message: "ERROR: The limit of 30 configured tunnel groups has been reached."

[Table 30-1](#) specifies the maximum VPN sessions and connection profiles for each ASA platform.

Table 30-1 Maximum VPN Sessions and Connection Profiles Per ASA Platform

	5505 Base/ Security Plus	5510/Base/ Security Plus	5520	5540	5550
Maximum VPN Sessions	10/25	250	750	5000	5000
Maximum Connection Profiles	15/30	255	755	5005	5005

Default IPsec Remote Access Tunnel Group Configuration

The contents of the default remote-access tunnel group are as follows:

```

tunnel-group DefaultRAGroup type ipsec-ra
tunnel-group DefaultRAGroup general-attributes
no address-pool
authentication-server-group LOCAL
no authorization-server-group
no accounting-server-group
default-group-policy DfltGrpPolicy
no dhcp-server
no nac-authentication-server-group
no strip-realm
no password-management
no override-account-disable
no strip-group
no authorization-required
authorization-dn-attributes CN OU
tunnel-group DefaultRAGroup ipsec-attributes
no pre-shared-key
peer-id-validate req
no chain
no trust-point
isakmp keepalive threshold 100 retry 2
isakmp ikev1-user-authentication xauth
tunnel-group DefaultRAGroup ppp-attributes
no authentication pap
authentication chap
authentication ms-chap-v1
no authentication ms-chap-v2
no authentication eap-proxy

```

Configuring IPsec Tunnel-Group General Attributes

The general attributes are common across more than one tunnel-group type. IPsec remote access and WebVPN tunnels share most of the same general attributes. IPsec LAN-to-LAN tunnels use a subset. Refer to the *Cisco Security Appliance Command Reference* for complete descriptions of all commands. The following sections describe, in order, how to configure IPsec remote-access tunnel groups, IPsec LAN-to-LAN tunnel groups, and WebVPN tunnel groups.

Configuring IPsec Remote-Access Tunnel Groups

Use an IPsec remote-access tunnel group when setting up a connection between a remote client and a central-site security appliance, using a hardware or software client. To configure an IPsec remote-access tunnel group, first configure the tunnel-group general attributes, then the IPsec remote-access attributes. An IPsec Remote Access VPN tunnel group applies only to remote-access IPsec client connections. To configure an IPsec remote-access tunnel group, see the following sections:

- [Specifying a Name and Type for the IPsec Remote Access Tunnel Group, page 30-6.](#)
- [Configuring IPsec Remote-Access Tunnel Group General Attributes, page 30-7.](#)
- [Configuring IPsec Remote-Access Tunnel Group IPsec Attributes, page 30-10.](#)

Specifying a Name and Type for the IPsec Remote Access Tunnel Group

Create the tunnel group, specifying its name and type, by entering the **tunnel-group** command. For an IPsec remote-access tunnel, the type is **ipsec-ra**

```
hostname(config)# tunnel-group tunnel_group_name type ipsec-ra
hostname(config)#
```

For example, to create an IPSec remote-access tunnel-group named TunnelGroup1, enter the following command:

```
hostname(config)# tunnel-group TunnelGroup1 type ipsec-ra
hostname(config)#
```

Configuring IPSec Remote-Access Tunnel Group General Attributes

To configure or change the tunnel group general attributes, specify the parameters in the following steps.

- Step 1** To configure the general attributes, enter tunnel-group general-attributes command, which enters tunnel-group general-attributes configuration mode. The prompt changes to indicate the change in mode.

```
hostname(config)# tunnel-group tunnel_group_name general-attributes
hostname(config-tunnel-general)#
```

- Step 2** Specify the name of the authentication-server group, if any, to use. If you want to use the LOCAL database for authentication if the specified server group fails, append the keyword **LOCAL**:

```
hostname(config-tunnel-general)# authentication-server-group [(interface_name)] groupname
[LOCAL]
hostname(config-tunnel-general)#
```

You can optionally configure interface-specific authentication by including the name of an interface after the group name. The interface name, which specifies where the IPSec tunnel terminates, must be enclosed in parentheses. The following command configures interface-specific authentication for the interface named test using the server named servergroup1 for authentication:

```
hostname(config-tunnel-general)# authentication-server-group (test) servergroup1
hostname(config-tunnel-general)#
```

- Step 3** Specify the name of the authorization-server group, if any, to use. When you configure this value, users must exist in the authorization database to connect:

```
hostname(config-tunnel-general)# authorization-server-group groupname
hostname(config-tunnel-general)#
```

For example, the following command specifies the use of the authorization-server group FinGroup:

```
hostname(config-tunnel-general)# authorization-server-group FinGroup
hostname(config-tunnel-general)#
```

- Step 4** Specify the name of the accounting-server group, if any, to use:

```
hostname(config-tunnel-general)# accounting-server-group groupname
hostname(config-tunnel-general)#
```

For example, the following command specifies the use of the accounting-server group named comptroller:

```
hostname(config-tunnel-general)# accounting-server-group comptroller
hostname(config-tunnel-general)#
```

- Step 5** Specify the name of the default group policy:

```
hostname(config-tunnel-general)# default-group-policy policyname
hostname(config-tunnel-general)#
```

The following example sets DfltGrpPolicy as the name of the default group policy:

```
hostname(config-tunnel-general)# default-group-policy DfltGrpPolicy
hostname(config-tunnel-general)#
```

- Step 6** Specify the names or IP addresses of the DHCP server (up to 10 servers), and the names of the DHCP address pools (up to 6 pools). The defaults are no DHCP server and no address pool.

```
hostname(config-tunnel-general)# dhcp-server server1 [...server10]
hostname(config-tunnel-general)# address-pool [(interface name)] address_pool1
[...address_pool6]
hostname(config-tunnel-general)#
```



Note The interface name must be enclosed in parentheses.

You configure address pools with the **ip local pool** command in global configuration mode.

- Step 7** Specify the name of the NAC authentication server group, if you are using Network Admission Control, to identify the group of authentication servers to be used for Network Admission Control posture validation. Configure at least one Access Control Server to support NAC. Use the **aaa-server** command to name the ACS group. Then use the **nac-authentication-server-group** command, using the same name for the server group.

The following example identifies acs-group1 as the authentication server group to be used for NAC posture validation:

```
hostname(config-group-policy)# nac-authentication-server-group acs-group1
hostname(config-group-policy)
```

The following example inherits the authentication server group from the default remote access group.

```
hostname(config-group-policy)# no nac-authentication-server-group
hostname(config-group-policy)
```



Note NAC requires a Cisco Trust Agent on the remote host.

- Step 8** Specify whether to strip the group or the realm from the username before passing it on to the AAA server. The default is not to strip either the group name or the realm.

```
hostname(config-tunnel-general)# strip-group
hostname(config-tunnel-general)# strip-realm
hostname(config-tunnel-general)#
```

A realm is an administrative domain. If you strip the realm, the security appliance uses the username and the group (if present) authentication. If you strip the group, the security appliance uses the username and the realm (if present) for authentication. Enter the **strip-realm** command to remove the realm qualifier, and use the **strip-group** command to remove the group qualifier from the username during authentication. If you remove both qualifiers, authentication is based on the *username* alone. Otherwise, authentication is based on the full *username@realm* or *username<delimiter> group* string. You must specify **strip-realm** if your server is unable to parse delimiters.

- Step 9** Optionally, if your server is a RADIUS, RADIUS with NT, or LDAP server, you can enable password management.

**Note**

If you are using an LDAP directory server for authentication, password management is supported with the Sun Microsystems JAVA System Directory Server (formerly named the Sun ONE Directory Server) and the Microsoft Active Directory.

- Sun—The DN configured on the security appliance to access a Sun directory server must be able to access the default password policy on that server. We recommend using the directory administrator, or a user with directory administrator privileges, as the DN. Alternatively, you can place an ACI on the default password policy.
- Microsoft—You must configure LDAP over SSL to enable password management with Microsoft Active Directory.

See the [“Setting the LDAP Server Type” section on page 13-7](#) for more information.

This feature, which is enabled by default, warns a user when the current password is about to expire. The default is to begin warning the user 14 days before expiration:

```
hostname(config-tunnel-general)# password-management
hostname(config-tunnel-general)#
```

If the server is an LDAP server, you can specify the number of days (0 through 180) before expiration to begin warning the user about the pending expiration:

```
hostname(config-tunnel-general)# password-management [password-expire in days n]
hostname(config-tunnel-general)#
```

**Note**

The **password-management** command, entered in tunnel-group general-attributes configuration mode replaces the deprecated **radius-with-expiry** command that was formerly entered in tunnel-group ipsec-attributes mode.

When you configure this command, the security appliance notifies the remote user at login that the user's current password is about to expire or has expired. The security appliance then offers the user the opportunity to change the password. If the current password has not yet expired, the user can still log in using that password. The security appliance ignores this command if RADIUS or LDAP authentication has not been configured.

Note that this does not change the number of days before the password expires, but rather, the number of days ahead of expiration that the security appliance starts warning the user that the password is about to expire.

If you do specify the **password-expire-in-days** keyword, you must also specify the number of days.

Specifying this command with the number of days set to 0 disables this command. The security appliance does not notify the user of the pending expiration, but the user can change the password after it expires.

See [Configuring Microsoft Active Directory Settings for Password Management, page 30-24](#) for more information.

**Note**

The **radius-with-expiry** command, formerly configured as part of tunnel-group ipsec-ra configuration, is deprecated. The **password-management** command, entered in tunnel-group general-attributes mode, replaces it.

- Step 10** Optionally, configure the ability to override an account-disabled indicator from a AAA server, by entering the **override-account-disable** command:

```
hostname(config-tunnel-general)# override-account-disable
hostname(config-tunnel-general)#
```

**Note**

Allowing `override-account-disable` is a potential security risk.

- Step 11** Specify the attribute or attributes to use in deriving a name for an authorization query from a certificate. This attribute specifies what part of the subject DN field to use as the username for authorization:

```
hostname(config-tunnel-general)# authorization-dn-attributes {primary-attribute
[secondary-attribute] | use-entire-name}
```

For example, the following command specifies the use of the CN attribute as the username for authorization:

```
hostname(config-tunnel-general)# authorization-dn-attributes CN
hostname(config-tunnel-general)#
```

The `authorization-dn-attributes` are **C** (Country), **CN** (Common Name), **DNQ** (DN qualifier), **EA** (E-mail Address), **GENQ** (Generational qualifier), **GN** (Given Name), **I** (Initials), **L** (Locality), **N** (Name), **O** (Organization), **OU** (Organizational Unit), **SER** (Serial Number), **SN** (Surname), **SP** (State/Province), **T** (Title), **UID** (User ID), and **UPN** (User Principal Name).

- Step 12** Specify whether to require a successful authorization before allowing a user to connect. The default is not to require authorization.

```
hostname(config-tunnel-general)# authorization-required
hostname(config-tunnel-general)#
```

Configuring IPsec Remote-Access Tunnel Group IPsec Attributes

To configure the IPsec attributes for a remote-access tunnel group, do the following steps. The following description assumes that you have already created the IPsec remote-access tunnel group. IPsec remote-access tunnel groups have more attributes than IPsec LAN-to-LAN tunnel groups:

- Step 1** To specify the attributes of an IPsec remote-access tunnel-group, enter `tunnel-group ipsec-attributes` mode by entering the following command. The prompt changes to indicate the mode change:

```
hostname(config)# tunnel-group tunnel-group-name ipsec-attributes
hostname(config-tunnel-ipsec)#
```

This command enters `tunnel-group ipsec-attributes` configuration mode, in which you configure the remote-access tunnel-group IPsec attributes.

For example, the following command designates that the `tunnel-group ipsec-attributes` mode commands that follow pertain to the tunnel group named TG1. Notice that the prompt changes to indicate that you are now in `tunnel-group ipsec-attributes` mode:

```
hostname(config)# tunnel-group TG1 type ipsec-ra
hostname(config)# tunnel-group TG1 ipsec-attributes
hostname(config-tunnel-ipsec)#
```

- Step 2** Specify the preshared key to support IKE connections based on preshared keys. For example, the following command specifies the preshared key `xyzx` to support IKE connections for an IPsec remote access tunnel group:

```
hostname(config-tunnel-ipsec)# pre-shared-key xyzx
```

```
hostname(config-tunnel-ipsec)#
```

- Step 3** Specify whether to validate the identity of the peer using the peer's certificate:

```
hostname(config-tunnel-ipsec)# peer-id-validate option
hostname(config-tunnel-ipsec)#
```

The available options are **req** (required), **cert** (if supported by certificate), and **nocheck** (do not check). The default is **req**.

For example, the following command specifies that peer-id validation is required:

```
hostname(config-tunnel-ipsec)# peer-id-validate req
hostname(config-tunnel-ipsec)#
```

- Step 4** Specify whether to

- Step 5** Specify whether to enable sending of a certificate chain. The following command includes the root certificate and any subordinate CA certificates in the transmission:

```
hostname(config-tunnel-ipsec)# chain
hostname(config-tunnel-ipsec)#
```

This attribute applies to all IPSec tunnel-group types.

- Step 6** Specify the name of a trustpoint that identifies the certificate to be sent to the IKE peer:

```
hostname(config-tunnel-ipsec)# trust-point trust-point-name
hostname(config-tunnel-ipsec)#
```

The following command specifies mytrustpoint as the name of the certificate to be sent to the IKE peer:

```
hostname(config-ipsec)# trust-point mytrustpoint
```

- Step 7** Specify the ISAKMP (IKE) keepalive threshold and the number of retries allowed.

```
hostname(config-tunnel-ipsec)# isakmp keepalive threshold <number> retry <number>
hostname(config-tunnel-ipsec)#
```

The **threshold** parameter specifies the number of seconds (10 through 3600) that the peer is allowed to idle before beginning keepalive monitoring. The **retry** parameter is the interval (2 through 10 seconds) between retries after a keepalive response has not been received. IKE keepalives are enabled by default. To disable IKE keepalives, enter the **no** form of the **isakmp** command:

For example, the following command sets the IKE keepalive threshold value to 15 seconds and sets the retry interval to 10 seconds:

```
hostname(config-tunnel-ipsec)# isakmp keepalive threshold 15 retry 10
hostname(config-tunnel-ipsec)#
```

The default value for the **threshold** parameter is 300 for remote-access and 10 for LAN-to-LAN, and the default value for the **retry** parameter is 2.

To specify that the central site ("head end") should never initiate ISAKMP monitoring, enter the following command:

```
hostname(config-tunnel-ipsec)# isakmp keepalive threshold infinite
hostname(config-tunnel-ipsec)#
```

- Step 8** Specify the ISAKMP hybrid authentication method, XAUTH or hybrid XAUTH.

You use **isakmp ikev1-user-authentication** command to implement hybrid XAUTH authentication when you need to use digital certificates for security appliance authentication and a different, legacy method for remote VPN user authentication, such as RADIUS, TACACS+ or SecurID. Hybrid XAUTH breaks phase 1 of IKE down into the following two steps, together called hybrid authentication:

- a. The security appliance authenticates to the remote VPN user with standard public key techniques. This establishes an IKE security association that is unidirectionally authenticated.
- b. An XAUTH exchange then authenticates the remote VPN user. This extended authentication can use one of the supported legacy authentication methods.



Note Before the authentication type can be set to hybrid, you must configure the authentication server, create a preshared key, and configure a trustpoint.

You can use the **isakmp ikev1-user-authentication** command with the optional **interface** parameter to specify a particular interface. When you omit the **interface** parameter, the command applies to all the interfaces and serves as a back-up when the per-interface command is not specified. When there are two **isakmp ikev1-user-authentication** commands specified for a tunnel group, and one uses the **interface** parameter and one does not, the one specifying the interface takes precedence for that particular interface.

For example, the following commands enable hybrid XAUTH on the inside interface for a tunnel group called example-group:

```
hostname(config)# tunnel-group example-group type ipsec-ra
hostname(config)# tunnel-group example-group ipsec-attributes
hostname(config-tunnel-ipsec)# isakmp ikev1-user-authentication (inside) hybrid
hostname(config-tunnel-ipsec)#
```

Configuring IPSec Remote-Access Tunnel Group PPP Attributes

To configure the Point-to-Point Protocol attributes for a remote-access tunnel group, do the following steps. PPP attributes apply *only* to IPSec remote-access tunnel groups. The following description assumes that you have already created the IPSec remote-access tunnel group.

- Step 1** Enter tunnel-group ppp-attributes configuration mode, in which you configure the remote-access tunnel-group PPP attributes, by entering the following command. The prompt changes to indicate the mode change:

```
hostname(config)# tunnel-group tunnel-group-name type ipsec-ra
hostname(config)# tunnel-group tunnel-group-name ppp-attributes
hostname(config-tunnel-ppp)#
```

For example, the following command designates that the tunnel-group ppp-attributes mode commands that follow pertain to the tunnel group named TG1. Notice that the prompt changes to indicate that you are now in tunnel-group ppp-attributes mode:

```
hostname(config)# tunnel-group TG1 type ipsec-ra
hostname(config)# tunnel-group TG1 ppp-attributes
hostname(config-tunnel-ppp)#
```

- Step 2** Specify whether to enable authentication using specific protocols for the PPP connection. The protocol value can be:
- pap—Enables the use of Password Authentication Protocol for the PPP connection.

- **chap**—Enables the use of Challenge Handshake Authentication Protocol for the PPP connection.
- **ms-chap-v1** or **ms-chap-v2**—Enables the use of Microsoft Challenge Handshake Authentication Protocol, version 1 or version 2 for the PPP connection.
- **eap**—Enables the use of Extensible Authentication protocol for the PPP connection.

CHAP and MSCHAPv1 are enabled by default.

The syntax of this command is:

```
hostname(config-tunnel-ppp) # authentication protocol
hostname(config-tunnel-ppp) #
```

To disable authentication for a specific protocol, use the **no** form of the command:

```
hostname(config-tunnel-ppp) # no authentication protocol
hostname(config-tunnel-ppp) #
```

For example, the following command enables the use of the PAP protocol for a PPP connection.

```
hostname(config-tunnel-ppp) # authentication pap
hostname(config-tunnel-ppp) #
```

The following command enables the use of the MS-CHAP, version 2 protocol for a PPP connection:

```
hostname(config-tunnel-ppp) # authentication ms-chap-v2
hostname(config-tunnel-ppp) #
```

The following command enables the use of the EAP-PROXY protocol for a PPP connection:

```
hostname(config-tunnel-ppp) # authentication pap
hostname(config-tunnel-ppp) #
```

The following command disables the use of the MS-CHAP, version 1 protocol for a PPP connection:

```
hostname(config-tunnel-ppp) # no authentication ms-chap-v1
hostname(config-tunnel-ppp) #
```

Configuring LAN-to-LAN Tunnel Groups

An IPsec LAN-to-LAN VPN tunnel group applies only to LAN-to-LAN IPsec client connections. While many of the parameters that you configure are the same as for IPsec remote-access tunnel groups, LAN-to-LAN tunnels have fewer parameters. To configure a LAN-to-LAN tunnel group, follow the steps in this section.

Default LAN-to-LAN Tunnel Group Configuration

The contents of the default LAN-to-LAN tunnel group are as follows:

```
tunnel-group DefaultL2LGroup type ipsec-l2l
tunnel-group DefaultL2LGroup general-attributes
  no accounting-server-group
  default-group-policy DfltGrpPolicy
tunnel-group DefaultL2LGroup ipsec-attributes
  no pre-shared-key
  peer-id-validate req
  no chain
  no trust-point
  isakmp keepalive threshold 10 retry 2
```

LAN-to-LAN tunnel groups have fewer parameters than remote-access tunnel groups, and most of these are the same for both groups. For your convenience in configuring the connection, they are listed separately here. Any parameters that you do not explicitly configure inherit their values from the default tunnel group.

Specifying a Name and Type for a LAN-to-LAN Tunnel Group

To specify a name and a type for a tunnel group, enter the **tunnel-group** command, as follows:

```
hostname(config)# tunnel-group tunnel_group_name type tunnel_type
```

For a LAN-to-LAN tunnel, the type is **ipsec-l2l**.; for example, to create the LAN-to-LAN tunnel group named docs, enter the following command:

```
hostname(config)# tunnel-group docs type ipsec-l2l
hostname(config)#
```

Configuring LAN-to-LAN Tunnel Group General Attributes

To configure the tunnel group general attributes, do the following steps:

-
- Step 1** Enter tunnel-group general-attributes mode by specifying the general-attributes keyword:

```
hostname(config)# tunnel-group tunnel-group-name general-attributes
hostname(config-tunnel-general)#
```

The prompt changes to indicate that you are now in config-general mode, in which you configure the tunnel-group general attributes.

For example, for the tunnel group named docs, enter the following command:

```
hostname(config)# tunnel-group docs general-attributes
hostname(config-tunnel-general)#
```

- Step 2** Specify the name of the accounting-server group, if any, to use:

```
hostname(config-tunnel-general)# accounting-server-group groupname
hostname(config-tunnel-general)#
```

For example, the following command specifies the use of the accounting-server group acctgserv1:

```
hostname(config-tunnel-general)# accounting-server-group acctgserv1
hostname(config-tunnel-general)#
```

- Step 3** Specify the name of the default group policy:

```
hostname(config-tunnel-general)# default-group-policy policyname
hostname(config-tunnel-general)#
```

For example, the following command specifies that the name of the default group policy is MyPolicy:

```
hostname(config-tunnel-general)# default-group-policy MyPolicy
hostname(config-tunnel-general)#
```

Configuring LAN-to-LAN IPsec Attributes

To configure the IPsec attributes, do the following steps:

- Step 1** To configure the tunnel-group IPsec attributes, enter tunnel-group ipsec-attributes configuration mode by entering the tunnel-group command with the IPsec-attributes keyword.

```
hostname(config)# tunnel-group tunnel-group-name ipsec-attributes
hostname(config-tunnel-ipsec)#
```

For example, the following command enters config-ipsec mode so you can configure the parameters for the tunnel group named TG1:

```
hostname(config)# tunnel-group TG1 ipsec-attributes
hostname(config-tunnel-ipsec)#
```

The prompt changes to indicate that you are now in tunnel-group ipsec-attributes configuration mode.

- Step 2** Specify the preshared key to support IKE connections based on preshared keys.

```
hostname(config-tunnel-ipsec)# pre-shared-key key
hostname(config-tunnel-ipsec)#
```

For example, the following command specifies the preshared key XYZX to support IKE connections for an IPsec LAN-to-LAN tunnel group:

```
hostname(config-tunnel-ipsec)# pre-shared-key xyzx
hostname(config-tunnel-general)#
```

- Step 3** Specify whether to validate the identity of the peer using the peer's certificate:

```
hostname(config-tunnel-ipsec)# peer-id-validate option
hostname(config-tunnel-ipsec)#
```

The available options are **req** (required), **cert** (if supported by certificate), and **nocheck** (do not check). The default is **req**. For example, the following command sets the peer-id-validate option to **nocheck**:

```
hostname(config-tunnel-ipsec)# peer-id-validate nocheck
hostname(config-tunnel-ipsec)#
```

- Step 4** Specify whether to enable sending of a certificate chain. This action includes the root certificate and any subordinate CA certificates in the transmission:

```
hostname(config-tunnel-ipsec)# chain
hostname(config-tunnel-ipsec)#
```

You can apply this attribute to all tunnel-group types.

- Step 5** Specify the name of a trustpoint that identifies the certificate to be sent to the IKE peer:

```
hostname(config-tunnel-ipsec)# trust-point trust-point-name
hostname(config-tunnel-ipsec)#
```

For example, the following command sets the trustpoint name to mytrustpoint:

```
hostname(config-tunnel-ipsec)# trust-point mytrustpoint
hostname(config-tunnel-ipsec)#
```

You can apply this attribute to all tunnel-group types.

- Step 6** Specify the ISAKMP(IKE) keepalive threshold and the number of retries allowed. The **threshold** parameter specifies the number of seconds (10 through 3600) that the peer is allowed to idle before beginning keepalive monitoring. The **retry** parameter is the interval (2 through 10 seconds) between retries after a keepalive response has not been received. IKE keepalives are enabled by default. To disable IKE keepalives, enter the **no** form of the **isakmp** command:

```
hostname(config)# isakmp keepalive threshold <number> retry <number>
hostname(config-tunnel-ipsec)#
```

For example, the following command sets the ISAKMP keepalive threshold to 15 seconds and sets the retry interval to 10 seconds.:

```
hostname(config-tunnel-ipsec)# isakmp keepalive threshold 15 retry 10
hostname(config-tunnel-ipsec)#
```

The default value for the **threshold** parameter for LAN-to-LAN is 10, and the default value for the retry parameter is 2.

To specify that the central site (“head end”) should never initiate ISAKMP monitoring, enter the following command:

```
hostname(config-tunnel-ipsec)# isakmp keepalive threshold infinite
hostname(config-tunnel-ipsec)#
```

- Step 7** Specify the ISAKMP hybrid authentication method, XAUTH or hybrid XAUTH.

You use **isakmp ikev1-user-authentication** command to implement hybrid XAUTH authentication when you need to use digital certificates for security appliance authentication and a different, legacy method for remote VPN user authentication, such as RADIUS, TACACS+ or SecurID. Hybrid XAUTH breaks phase 1 of IKE down into the following two steps, together called hybrid authentication:

- a. The security appliance authenticates to the remote VPN user with standard public key techniques. This establishes an IKE security association that is unidirectionally authenticated.
- b. An XAUTH exchange then authenticates the remote VPN user. This extended authentication can use one of the supported legacy authentication methods.



Note Before the authentication type can be set to hybrid, you must configure the authentication server, create a preshared key, and configure a trustpoint.

You can use the **isakmp ikev1-user-authentication** command with the optional **interface** parameter to specify a particular interface. When you omit the **interface** parameter, the command applies to all the interfaces and serves as a back-up when the per-interface command is not specified. When there are two **isakmp ikev1-user-authentication** commands specified for a tunnel group, and one uses the **interface** parameter and one does not, the one specifying the interface takes precedence for that particular interface.

For example, the following commands enable hybrid XAUTH on the inside interface for a tunnel group called example-group:

```
hostname(config)# tunnel-group example-group type ipsec-ra
hostname(config)# tunnel-group example-group ipsec-attributes
hostname(config-tunnel-ipsec)# isakmp ikev1-user-authentication (inside) hybrid
hostname(config-tunnel-ipsec)#
```

Configuring WebVPN Tunnel Groups

A WebVPN tunnel group applies only to WebVPN connections. The tunnel-group general attributes for WebVPN tunnel groups are the same as those of IPSec remote-access tunnel groups, except that the tunnel-group type is **webvpn** and the **strip-group** and **strip-realm** commands do not apply. You define the WebVPN-specific attributes separately. The following sections describe how to configure WebVPN tunnel groups.

Specifying a Name and Type for a WebVPN Tunnel Group

Create the tunnel group, specifying its name and type by entering the **tunnel-group** command in global configuration mode. For an IPSec remote-access tunnel, the type is **webvpn**.

```
hostname(config)# tunnel-group tunnel_group_name type webvpn
hostname(config)#
```

For example, to create a WebVPN tunnel-group named TunnelGroup3, enter the following command:

```
hostname(config)# tunnel-group TunnelGroup3 type webvpn
hostname(config)#
```

Configuring WebVPN Tunnel-Group General Attributes

To configure or change the tunnel group general attributes, specify the parameters in the following steps.

-
- Step 1** To configure the general attributes, enter **tunnel-group general-attributes** command, which enters tunnel-group general-attributes configuration mode. Note that the prompt changes:

```
hostname(config)# tunnel-group tunnel_group_name general-attributes
hostname(config-tunnel-general)#
```

To configure the general attributes for TunnelGroup3, created in the previous section, enter the following command:

```
hostname(config)# tunnel-group TunnelGroup3 general-attributes
hostname(config-tunnel-general)#
```

- Step 2** Specify the name of the authentication-server group, if any, to use. If you want to use the LOCAL database for authentication if the specified server group fails, append the keyword **LOCAL**:

```
hostname(config-tunnel-general)# authentication-server-group groupname [LOCAL]
hostname(config-tunnel-general)#
```

For example, to configure the authentication server group named test, and to provide fallback to the LOCAL server if the authentication server group fails, enter the following command:

```
hostname(config-tunnel-general)# authentication-server-group test LOCAL
hostname(config-tunnel-general)#
```

The authentication-server-group name identifies a previously configured authentication server or group of servers. Use the **aaa-server** command to configure authentication servers. The maximum length of the group tag is 16 characters.

You can also configure interface-specific authentication by including the name of an interface in parentheses before the group name. The following interfaces are available by default:

- inside—Name of interface GigabitEthernet0/1

- **outside**— Name of interface GigabitEthernet0/0

Other interfaces you have configured (using the **interface** command) are also available. The following command configures interface-specific authentication for the interface named **outside** using the server **servergroup1** for authentication:

```
hostname(config-tunnel-general)# authentication-server-group (outside) servergroup1
hostname(config-tunnel-general)#
```

- Step 3** Optionally, specify the name of the authorization-server group, if any, to use. If you are not using authorization, go to Step 6. When you configure this value, users must exist in the authorization database to connect:

```
hostname(config-tunnel-general)# authorization-server-group groupname
hostname(config-tunnel-general)#
```

Use the **aaa-server** command to configure authorization servers. The maximum length of the group tag is 16 characters.

For example, the following command specifies the use of the authorization-server group **FinGroup**:

```
hostname(config-tunnel-general)# authorization-server-group FinGroup
hostname(config-tunnel-general)#
```

- Step 4** Specify whether to require a successful authorization before allowing a user to connect. The default is not to require authorization.

```
hostname(config-tunnel-general)# authorization-required
hostname(config-tunnel-general)#
```

- Step 5** Specify the attribute or attributes to use in deriving a name for an authorization query from a certificate. This attribute specifies what part of the subject DN field to use as the username for authorization:

```
hostname(config-tunnel-general)# authorization-dn-attributes {primary-attribute
[secondary-attribute] | use-entire-name}
```

For example, the following command specifies the use of the **CN** attribute as the username for authorization:

```
hostname(config-tunnel-general)# authorization-dn-attributes CN
hostname(config-tunnel-general)#
```

The authorization-dn-attributes are **C** (Country), **CN** (Common Name), **DNQ** (DN qualifier), **EA** (E-mail Address), **GENQ** (Generational qualifier), **GN** (Given Name), **I** (Initials), **L** (Locality), **N** (Name), **O** (Organization), **OU** (Organizational Unit), **SER** (Serial Number), **SN** (Surname), **SP** (State/Province), **T** (Title), **UID** (User ID), and **UPN** (User Principal Name).

- Step 6** Optionally, specify the name of the accounting-server group, if any, to use. If you are not using accounting, go to Step 7. Use the **aaa-server** command to configure accounting servers. The maximum length of the group tag is 16 characters.:

```
hostname(config-tunnel-general)# accounting-server-group groupname
hostname(config-tunnel-general)#
```

For example, the following command specifies the use of the accounting-server group **comptroller**:

```
hostname(config-tunnel-general)# accounting-server-group comptroller
hostname(config-tunnel-general)#
```

- Step 7** Optionally, specify the name of the default group policy. The default value is **DfltGrpPolicy**:

```
hostname(config-tunnel-general)# default-group-policy policyname
hostname(config-tunnel-general)#
```

The following example sets MyDfltGrpPolicy as the name of the default group policy:

```
hostname(config-tunnel-general)# default-group-policy MyDfltGrpPolicy
hostname(config-tunnel-general)#
```

- Step 8** Optionally, specify the name or IP address of the DHCP server (up to 10 servers), and the names of the DHCP address pools (up to 6 pools). Separate the list items with spaces. The defaults are no DHCP server and no address pool.

```
hostname(config-tunnel-general)# dhcp-server server1 [...server10]
hostname(config-tunnel-general)# address-pool [(interface name)] address_pool1
[...address_pool6]
hostname(config-tunnel-general)#
```



Note The interface name must be enclosed in parentheses.

You configure address pools with the **ip local pool** command in global configuration mode. See [Chapter 31, “Configuring IP Addresses for VPNs”](#) for information about configuring address pools.

- Step 9** Optionally, if your server is a RADIUS, RADIUS with NT, or LDAP server, you can enable password management.



Note If you are using an LDAP directory server for authentication, password management is supported with the Sun Microsystems JAVA System Directory Server (formerly named the Sun ONE Directory Server) and the Microsoft Active Directory.

- Sun—The DN configured on the security appliance to access a Sun directory server must be able to access the default password policy on that server. We recommend using the directory administrator, or a user with directory administrator privileges, as the DN. Alternatively, you can place an ACI on the default password policy.
- Microsoft—You must configure LDAP over SSL to enable password management with Microsoft Active Directory.

See the [“Setting the LDAP Server Type”](#) section on page 13-7 for more information.

This feature, which is enabled by default, warns a user when the current password is about to expire. The default is to begin warning the user 14 days before expiration:

```
hostname(config-tunnel-general)# password-management
hostname(config-tunnel-general)#
```

If the server is an LDAP server, you can specify the number of days (0 through 180) before expiration to begin warning the user about the pending expiration:

```
hostname(config-tunnel-general)# password-management [password-expire in days n]
hostname(config-tunnel-general)#
```



Note The **password-management** command, entered in tunnel-group general-attributes configuration mode replaces the deprecated **radius-with-expiry** command that was formerly entered in tunnel-group ipsec-attributes mode.

When you configure this command, the security appliance notifies the remote user at login that the user's current password is about to expire or has expired. The security appliance then offers the user the opportunity to change the password. If the current password has not yet expired, the user can still log in using that password. The security appliance ignores this command if RADIUS or LDAP authentication has not been configured.

Note that this does not change the number of days before the password expires, but rather, the number of days ahead of expiration that the security appliance starts warning the user that the password is about to expire.

If you do specify the **password-expire-in-days** keyword, you must also specify the number of days.

See [Configuring Microsoft Active Directory Settings for Password Management, page 30-24](#) for more information.

- Step 10** Specifying this command with the number of days set to 0 disables this command. The security appliance does not notify the user of the pending expiration, but the user can change the password after it expires. Optionally, configure the ability to override an account-disabled indicator from the AAA server, by entering the **override-account-disable** command:

```
hostname(config-tunnel-general)# override-account-disable
hostname(config-tunnel-general)#
```



Note

Allowing override account-disabled is a potential security risk.

Configuring WebVPN Tunnel-Group WebVPN Attributes

To configure the parameters specific to a WebVPN tunnel group, follow the steps in this section.

- Step 1** To specify the attributes of a WebVPN tunnel-group, enter tunnel-group webvpn-attributes mode by entering the following command. The prompt changes to indicate the mode change:

```
hostname(config)# tunnel-group tunnel-group-name webvpn-attributes
hostname(config-tunnel-ipsec)#
```

For example, to specify the webvpn-attributes for the WebVPN tunnel-group named sales, enter the following command:

```
hostname(config)# tunnel-group sales webvpn-attributes
hostname(config-tunnel-webvpn)#
```

- Step 2** To specify the authentication method to use: AAA, digital certificates, or both, enter the **authentication** command. You can specify either aaa or certificate or both, in any order.

```
hostname(config-tunnel-webvpn)# authentication authentication_method
hostname(config-tunnel-webvpn)#
```

For example, The following command allows both AAA and certificate authentication:

```
hostname(config-tunnel-webvpn)# authentication aaa certificate
hostname(config-tunnel-webvpn)#
```


Applying Customization

Customizations determine the appearance of the windows that the user sees upon login. You configure the customization parameters as part of configuring WebVPN.

To apply a previously defined web-page customization to change the look-and-feel of the web page that the user sees at login, enter the customization command in username webvpn configuration mode:

```
hostname(config-username-webvpn)# customization {none | value customization_name}
hostname(config-username-webvpn)#
```

For example, to use the customization named blueborder, enter the following command:

```
hostname(config-username-webvpn)# customization value blueborder
hostname(config-username-webvpn)#
```

You configure the customization itself by entering the **customization** command in WebVPN mode.

The following example shows a command sequence that first establishes a WebVPN customization named “123” that defines a password prompt. The example then defines a WebVPN tunnel-group named “test” and uses the **customization** command to specifies the use of the WebVPN customization named “123”:

```
hostname(config)# webvpn
hostname(config-webvpn)# customization 123
hostname(config-webvpn-custom)# password-prompt Enter password
hostname(config-webvpn)# exit
hostname(config)# tunnel-group test type webvpn
hostname(config)# tunnel-group test webvpn-attributes
hostname(config-tunnel-webvpn)# customization value 123
hostname(config-tunnel-webvpn)#
```

Step 3

The security appliance queries NetBIOS name servers to map NetBIOS names to IP addresses. WebVPN requires NetBIOS to access or share files on remote systems. WebVPN uses NetBIOS and the CIFS protocol to access or share files on remote systems. When you attempt a file-sharing connection to a Windows computer by using its computer name, the file server you specify corresponds to a specific NetBIOS name that identifies a resource on the network.

To make the NBNS function operational, you must configure at least one NetBIOS server (host). You can configure up to three NBNS servers for redundancy. The security appliance uses the first server on the list for NetBIOS/CIFS name resolution. If the query fails, it uses the next server.

To specify the name of the NBNS (NetBIOS Name Service) server to use for CIFS name resolution, use the **nbns-server** command. You can enter up to three server entries. The first server you configure is the primary server, and the others are backups, for redundancy. You can also specify whether this is a master browser (rather than just a WINS server), the timeout interval, and the number of retries. A WINS server or a master browser is typically on the same network as the security appliance, or reachable from that network. You must specify the timeout interval before the number of retries:

```
hostname(config-tunnel-webvpn)# nbns-server {host-name | IP_address} [master]
[timeout seconds] [retry number]
hostname(config-tunnel-webvpn)#
```

For example, to configure the server named nbnsprimary as the primary server and the server 192.168.2.2 as the secondary server, each allowing three retries and having a 5-second timeout, enter the following command:

```
hostname(config)# name 192.168.2.1 nbnsprimary
hostname(config-tunnel-webvpn)# nbns-server nbnsprimary master timeout 5 retry 3
hostname(config-tunnel-webvpn)# nbns-server 192.168.2.2 timeout 5 retry 3
hostname(config-tunnel-webvpn)#
```

The timeout interval can range from 1 through 30 seconds (default 2), and the number of retries can be in the range 0 through 10 (default 2).

The **nbns-server** command in tunnel-group webvpn-attributes configuration mode replaces the deprecated **nbns-server** command in webvpn configuration mode.

- Step 4** To specify alternative names for the group, use the **group-alias** command. Specifying the group alias creates one or more alternate names by which the user can refer to a tunnel-group. The group alias that you specify here appears in the drop-down list on the user's login page. Each group can have multiple aliases or no alias, each specified in separate commands. This feature is useful when the same group is known by several common names, such as "Devtest" and "QA".

For each group alias, enter a **group-alias** command. Each alias is enabled by default. You can optionally explicitly enable or disable each alias:

```
hostname(config-tunnel-webvpn)# group-alias alias [enable | disable]
hostname(config-tunnel-webvpn)#
```

For example, to enable the aliases QA and Devtest for a tunnel-group named QA, enter the following commands:

```
hostname(config-tunnel-webvpn)# group-alias QA enable
hostname(config-tunnel-webvpn)# group-alias Devtest enable
hostname(config-tunnel-webvpn)#
```



Note

The WebVPN tunnel-group-list must be enabled for the (dropdown) group list to appear.

- Step 5** To specify incoming URLs or IP addresses for the group, use the **group-url** command. Specifying a group URL or IP address eliminates the need for the user to select a group at login. When a user logs in, the security appliance looks for the user's incoming URL or address in the tunnel-group-policy table. If it finds the URL or address and if group-url is enabled in the tunnel group, then the security appliance automatically selects the associated tunnel group and presents the user with only the username and password fields in the login window. This simplifies the user interface and has the added advantage of never exposing the list of groups to the user. The login window that the user sees uses the customizations configured for that tunnel group.

If the URL or address is disabled and group-alias is configured, then the dropdown list of groups is also displayed, and the user must make a selection.

You can configure multiple URLs or addresses (or none) for a group. Each URL or address can be enabled or disabled individually. You must use a separate **group-url** command for each URL or address specified. You must specify the entire URL or address, including either the http or https protocol.

You cannot associate the same URL or address with multiple groups. The security appliance verifies the uniqueness of the URL or address before accepting the URL or address for a tunnel group.

For each group URL or address, enter a **group-URL** command. You can optionally explicitly enable (the default) or disable each URL or alias:

```
hostname(config-tunnel-webvpn)# group-url url [enable | disable]
hostname(config-tunnel-webvpn)#
```

For example, to enable the group URLs http://www.cisco.com and http://192.168.10.10 for the tunnel-group named RadiusServer, enter the following commands:

```
hostname(config)# tunnel-group RadiusServer type webvpn
hostname(config)# tunnel-group RadiusServer general-attributes
hostname(config-tunnel-general)# authentication server-group RADIUS
hostname(config-tunnel-general)# accounting-server-group RADIUS
hostname(config-tunnel-general)# tunnel-group RadiusServer webvpn-attributes
```

```
hostname(config-tunnel-webvpn) # group-alias "Cisco Remote Access" enable
hostname(config-tunnel-webvpn) # group-url http://www.cisco.com enable
hostname(config-tunnel-webvpn) # group-url http://192.168.10.10 enable
hostname(config-tunnel-webvpn) #
```

For a more extensive example, see [Customizing Login Windows for WebVPN Users](#), page 30-23.

- Step 6** To specify the DNS server to use for a WebVPN tunnel group, enter the **dns-group** command. The default value is DefaultDNS:

```
hostname(config-tunnel-webvpn) # dns-group {hostname | ip_address}
hostname(config-tunnel-webvpn) #
```

The **dns-group** command resolves the hostname to the appropriate DNS server for the tunnel group. For example, to specify the use of the DNS server named server1, enter the following command:

```
hostname(config) # name 10.10.10.1 server1
hostname(config-tunnel-webvpn) # dns-group server1
hostname(config-tunnel-webvpn) #
```

- Step 7** (Optional) To specify a VPN feature policy if you use the Cisco Secure Desktop Manager to set the Group-Based Policy attribute to “Use Failure Group-Policy” or “Use Success Group-Policy, if criteria match,” use the **hic-fail-group-policy** command. The default value is DfltGrpPolicy.

```
hostname(config-tunnel-webvpn) # hic-fail-group-policy name
hostname(config-tunnel-webvpn) #
```

Name is the name of a group policy created for a WebVPN tunnel group.

This policy is an alternative group policy to differentiate access rights for the following CSD clients:

- Clients that match a CSD location entry set to “Use Failure Group-Policy.”
- Clients that match a CSD location entry set to “Use Success Group-Policy, if criteria match,” and then fail to match the configured Group-Based Policy criteria. For more information, see the *Cisco Secure Desktop Configuration Guide for Cisco ASA 5500 Series Administrators*.

The following example specifies an alternative group policy named group2:

```
hostname(config-tunnel-webvpn) # hic-fail-group-policy group2
hostname(config-tunnel-webvpn) #
```



Note The security appliance does not use this attribute if you set the VPN feature policy to “Always use Success Group-Policy.”

For more information, see the *Cisco Secure Desktop Configuration Guide for Cisco ASA 5500 Series Administration Guide*.

Customizing Login Windows for WebVPN Users

You can set up different login windows for different groups by using a combination of customization profiles and tunnel groups. For example, assuming that you had created a customization profile called salesgui, you can create a WebVPN tunnel group called sales that uses that customization profile, as the following example shows:

- Step 1** In webvpn mode, define a WebVPN customization, in this case named salesgui and change the default logo to mycompanylogo.gif. You must have previously loaded mycompanylogo.gif onto the flash memory of the security appliance and saved the configuration. See the WebVPN chapter for details.

```
hostname# webvpn
hostname (config-webvpn)# customization value salesgui
hostname (config-webvpn-custom)# logo file disk0:\mycompanylogo.gif
hostname (config-webvpn-custom)#
```

- Step 2** In global configuration mode, set up a username and associate with it the WebVPN customization you've just defined:

```
hostname# username seller attributes
hostname (config-username)# webvpn
hostname (config-username-webvpn)# customization value salesgui
hostname (config-username-webvpn)# exit
hostname (config-username)# exit
hostname#
```

- Step 3** In global configuration mode, create a WebVPN tunnel-group named sales:

```
hostname# tunnel-group sales type webvpn
hostname (config-tunnel-webvpn)#
```

- Step 4** Specify that you want to use the salesgui customization for this tunnel group:

```
hostname# tunnel-group sales webvpn-attributes
hostname (config-tunnel-webvpn)# customization salesgui
```

- Step 5** Set the group URL to the address that the user enters into the browser to log in to the security appliance; for example, if the security appliance has the IP address 192.168.3.3, set the group URL to https://192.168.3.3:

```
hostname (config-tunnel-webvpn)# group-url https://192.168.3.3.
hostname (config-tunnel-webvpn)#
```

If a port number is required for a successful login, include the port number, preceded by a colon. The security appliance maps this URL to the sales tunnel group and applies the salesgui customization profile to the login screen that the user sees upon logging in to https://192.168.3.3.

Configuring Microsoft Active Directory Settings for Password Management



Note

If you are using an LDAP directory server for authentication, password management is supported with the Sun Microsystems JAVA System Directory Server (formerly named the Sun ONE Directory Server) and the Microsoft Active Directory.

- Sun—The DN configured on the security appliance to access a Sun directory server must be able to access the default password policy on that server. We recommend using the directory administrator, or a user with directory administrator privileges, as the DN. Alternatively, you can place an ACI on the default password policy.
- Microsoft—You must configure LDAP over SSL to enable password management with Microsoft Active Directory.

See the [“Setting the LDAP Server Type”](#) section on page 13-7 for more information.

To use password management with Microsoft Active Directory, you must set certain Active Directory parameters as well as configuring password management on the security appliance. This section describes the Active Directory settings associated with various password management actions. These descriptions assume that you have also enabled password management on the security appliance and configured the corresponding password management attributes. The specific steps in the following sections refer to Active Directory terminology under Windows 2000.


- [Using Active Directory to Force the User to Change Password at Next Logon, page 30-25.](#)
- [Using Active Directory to Specify Maximum Password Age, page 30-27.](#)
- [Using Active Directory to Override an Account Disabled AAA Indicator, page 30-28](#)
- [Using Active Directory to Enforce Password Complexity, page 30-30.](#)

The following sections assume that you are using an LDAP directory server for authentication.

Using Active Directory to Force the User to Change Password at Next Logon


To force a user to change the user password at the next logon, specify the **password-management** command in tunnel-group general-attributes configuration mode on the security appliance and do the following steps under Active Directory:

-
- Step 1** Select to Start > Programs > Administrative Tools > Active Directory Users and Computers ([Figure 30-1](#)).

Figure 30-1 Active Directory—Administrative Tools Menu

Step 2 Right-click Username > Properties > Account.

Step 3 Check the check box for User must change password at next logon ([Figure 30-2](#)).


Figure 30-2 Active Directory—User Must Change Password at Next Logon

The next time this user logs on, the security appliance displays the following prompt: “New password required. Password change required. You must enter a new password with a minimum length n to continue.” You can set the minimum required password length, n , as part of the Active Directory configuration at Start > Programs > Administrative Tools > Domain Security Policy > Windows Settings > Security Settings > Account Policies > Password Policy. Select Minimum password length.

Using Active Directory to Specify Maximum Password Age

To enhance security, you can specify that passwords expire after a certain number of days. To specify a maximum password age for a user password, specify the **password-management** command in tunnel-group general-attributes configuration mode on the security appliance and do the following steps under Active Directory:

- Step 1** Select Start > Programs > Administrative Tools > Domain Security Policy > Windows Settings > Security Settings > Account Policies > Password Policy.
- Step 2** Double-click Maximum password age. This opens the Security Policy Setting dialog box.
- Step 3** Check the Define this policy setting check box and specify the maximum password age, in days, that you want to allow.

Figure 30-3 Active Directory—Maximum Password Age**Note**

The **radius-with-expiry** command, formerly configured as part of tunnel-group ipsec-ra configuration to perform the password age function, is deprecated. The **password-management** command, entered in tunnel-group general-attributes mode, replaces it.


Using Active Directory to Override an Account Disabled AAA Indicator

To override an account-disabled indication from a AAA server, specify the **override-account-disable** command in tunnel-group general-attributes configuration mode on the security appliance and do the following steps under Active Directory:

**Note**

Allowing override account-disabled is a potential security risk.

- Step 1** Select Start > Programs > Administrative Tools > Active Directory Users and Computers.
- Step 2** Right-click Username > Properties > Account and select Disable Account from the menu.


Figure 30-4 Active Directory—Override Account Disabled

The user should be able to log on successfully, even though a AAA server provides an account-disabled indicator.

Using Active Directory to Enforce Minimum Password Length

To enforce a minimum length for passwords, specify the **password-management** command in tunnel-group general-attributes configuration mode on the security appliance and do the following steps under Active Directory:


- Step 1** Select Start > Programs > Administrative Tools > Domain Security Policy.
- Step 2** Select Windows Settings > Security Settings > Account Policies > Password Policy.
- Step 3** Double-click Minimum Password Length. This opens the Security Policy Setting dialog box.
- Step 4** Check the Define this policy setting check box and specify the minimum number of characters that the password must contain.

Figure 30-5 Active Directory—Minimum Password Length

Using Active Directory to Enforce Password Complexity

To enforce complex passwords—for example, to require that a password contain upper- and lowercase letters, numbers, and special characters—specify the **password-management** command in tunnel-group general-attributes configuration mode on the security appliance and do the following steps under Active Directory:

- Step 1** Select Start > Programs > Administrative Tools > Domain Security Policy. Select Windows Settings > Security Settings > Account Policies > Password Policy.
- Step 2** Double-click Password must meet complexity requirements to open the Security Policy Setting dialog box.
- Step 3** Check the Define this policy setting check box and select Enable.

Figure 30-6 Active Directory—Enforce Password Complexity

Enforcing password complexity takes effect only when the user changes passwords; for example, when you have configured Enforce password change at next login or Password expires in n days. At login, the user receives a prompt to enter a new password, and the system will accept only a complex password.

Group Policies

This section describes group policies and how to configure them. It includes the following sections:

- [Default Group Policy, page 30-32](#)
- [Configuring Group Policies, page 30-34](#)

A group policy is a set of user-oriented attribute/value pairs for IPSec connections that are stored either internally (locally) on the device or externally on a RADIUS server. The tunnel group uses a group policy that sets terms for user connections after the tunnel is established. Group policies let you apply whole sets of attributes to a user or a group of users, rather than having to specify each attribute individually for each user.

Enter the **group-policy** commands in global configuration mode to assign a group policy to users or to modify a group policy for specific users.

The security appliance includes a default group policy. In addition to the default group policy, which you can modify but not delete, you can create one or more group policies specific to your environment.

You can configure internal and external group policies. Internal groups are configured on the security appliance's internal database. External groups are configured on an external authentication server, such as RADIUS. Group policies include the following attributes:

- Identity
- Server definitions

- Client firewall settings
- Tunneling protocols
- IPSec settings
- Hardware client settings
- Filters
- Client configuration settings
- WebVPN functions
- Connection settings

Default Group Policy

The security appliance supplies a default group policy. You can modify this default group policy, but you cannot delete it. A default group policy, named `DfltGrpPolicy`, always exists on the security appliance, but this default group policy does not take effect unless you configure the security appliance to use it. When you configure other group policies, any attribute that you do not explicitly specify takes its value from the default group policy. To view the default group policy, enter the following command:

```
hostname(config)# show running-config all group-policy DfltGrpPolicy
hostname(config)#
```

To configure the default group policy, enter the following command:

```
hostname(config)# group-policy DfltGrpPolicy internal
hostname(config)#
```



Note

The default group policy is always internal. Despite the fact that the command syntax is

```
hostname(config)# group-policy DfltGrpPolicy {internal | external}, you cannot change the type to external.
```

To change any of the attributes of the default group policy, use the **group-policy attributes** command to enter attributes mode, then specify the commands to change whatever attributes that you want to modify:

```
hostname(config)# group-policy DfltGrpPolicy attributes
```



Note

The attributes mode applies only to internal group policies.

The default group policy, `DfltGrpPolicy`, that the security appliance provides is as follows:

```
group-policy DfltGrpPolicy internal
group-policy DfltGrpPolicy attributes
  banner none
  wins-server none
  dns-server none
  dhcp-network-scope none
  vpn-access-hours none
  vpn-simultaneous-logins 2000
  vpn-idle-timeout none
  vpn-session-timeout none
  vpn-filter none
```

```

vpn-tunnel-protocol IPSec webvpn
password-storage enable
ip-comp disable
re-xauth disable
group-lock none
pfs disable
ipsec-udp disable
ipsec-udp-port 10000
split-tunnel-policy tunnelall
split-tunnel-network-list none
default-domain none
split-dns none
intercept-dhcp 255.255.255.255 disable
secure-unit-authentication disable
user-authentication disable
user-authentication-idle-timeout 30
ip-phone-bypass disable
leap-bypass disable
nem disable
backup-servers keep-client-config
msie-proxy server none
msie-proxy method no-modify
msie-proxy except-list none
msie-proxy local-bypass disable
nac disable
nac-sq-period 300
nac-reval-period 36000
nac-default-acl none
address-pools value vpn_users
client-firewall none
client-access-rule none
webvpn
  functions url-entry file-access file-browsing
  html-content-filter none
  homepage none
  keep-alive-ignore 4
  http-comp gzip
  filter none
  url-list value MyURLs
  customization value DfltCustomization
  port-forward none
  port-forward-name value Application Access
  sso-server none
  deny-message value Login was successful, but because certain criteria have not been met
  or due to some specific group policy, you do not have permission to use any of the VPN
  features. Contact your IT administrator for more information
  svc none
  svc keep-installer none
  svc keepalive none
  svc rekey time none
  svc rekey method none
  svc dpd-interval client none
  svc dpd-interval gateway none
  svc compression deflate
  no vpn-nac-exempt
hostname(config-group-policy)#

```

You can modify the default group policy, and you can also create one or more group policies specific to your environment.

Configuring Group Policies

A group policy can apply to any kind of tunnel. In each case, if you do not explicitly define a parameter, the group takes the value from the default group policy. To configure a group policy, follow the steps in the subsequent sections.

Configuring an External Group Policy

External group policies take their attribute values from the external server that you specify. For an external group policy, you must identify the AAA server group that the security appliance can query for attributes and specify the password to use when retrieving attributes from the external AAA server group. If you are using an external authentication server, and if your external group-policy attributes exist in the same RADIUS server as the users that you plan to authenticate, you have to make sure that there is no name duplication between them.

**Note**

External group names on the security appliance refer to user names on the RADIUS server. In other words, if you configure external group X on the security appliance, the RADIUS server sees the query as an authentication request for user X. So external groups are really just user accounts on the RADIUS server that have special meaning to the security appliance. If your external group attributes exist in the same RADIUS server as the users that you plan to authenticate, there must be no name duplication between them.

The security appliance supports user authorization on an external LDAP or RADIUS server. Before you configure the security appliance to use an external server, you must configure the server with the correct security appliance authorization attributes and, from a subset of these attributes, assign specific permissions to individual users. Follow the instructions in [Appendix E, “Configuring an External Server for Authorization and Authentication”](#) to configure your external server.

To configure an external group policy, do the following steps specify a name and type for the group policy, along with the server-group name and a password:

```
hostname(config)# group-policy group_policy_name type server-group server_group_name  
password server_password  
hostname(config)#
```

**Note**

For an external group policy, RADIUS is the only supported AAA server type.

For example, the following command creates an external group policy named ExtGroup that gets its attributes from an external RADIUS server named ExtRAD and specifies that the password to use when retrieving the attributes is newpassword:

```
hostname(config)# group-policy ExtGroup external server-group ExtRAD password newpassword  
hostname(config)#
```

**Note**

You can configure several vendor-specific attributes (VSAs), as described in [Appendix E, “Configuring an External Server for Authorization and Authentication”](#). If a RADIUS server is configured to return the Class attribute (#25), the security appliance uses that attribute to authenticate the Group Name. On the RADIUS server, the attribute must be formatted as: OU=groupname; where groupname is identical to the Group Name configured on the security appliance—for example, OU=Finance.

Configuring an Internal Group Policy

To configure an internal group policy, specify a name and type for the group policy:

```
hostname(config)# group-policy group_policy_name type
hostname(config)#
```

For example, the following command creates the internal group policy named GroupPolicy1:

```
hostname(config)# group-policy GroupPolicy1 internal
hostname(config)#
```

The default type is **internal**.

You can initialize the attributes of an internal group policy to the values of a preexisting group policy by appending the keyword **from** and specifying the name of the existing policy:

```
hostname(config)# group-policy group_policy_name internal from group_policy_name
hostname(config-group-policy)#
hostname(config-group-policy)#
```

Configuring Group Policy Attributes

For internal group policies, you can specify particular attribute values. To begin, enter group-policy attributes mode, by entering the **group-policy attributes** command in global configuration mode.

```
hostname(config)# group-policy name attributes
hostname(config-group-policy)#
```

The prompt changes to indicate the mode change. The group-policy-attributes mode lets you configure attribute-value pairs for a specified group policy. In group-policy-attributes mode, explicitly configure the attribute-value pairs that you do not want to inherit from the default group. The commands to do this are described in the following sections.

Configuring WINS and DNS Servers

You can specify primary and secondary WINS servers and DNS servers. The default value in each case is none. To specify these servers, do the following steps:

Step 1 Specify the primary and secondary WINS servers:

```
hostname(config-group-policy)# wins-server value {ip_address [ip_address] | none}
hostname(config-group-policy)#
```

The first IP address specified is that of the primary WINS server. The second (optional) IP address is that of the secondary WINS server. Specifying the **none** keyword instead of an IP address sets WINS servers to a null value, which allows no WINS servers and prevents inheriting a value from a default or specified group policy.

Every time that you enter the **wins-server** command, you overwrite the existing setting. For example, if you configure WINS server x.x.x.x and then configure WINS server y.y.y.y, the second command overwrites the first, and y.y.y.y becomes the sole WINS server. The same is true for multiple servers. To add a WINS server rather than overwrite previously configured servers, include the IP addresses of all WINS servers when you enter this command.

The following example shows how to configure WINS servers with the IP addresses 10.10.10.15 and 10.10.10.30 for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# wins-server value 10.10.10.15 10.10.10.30
hostname(config-group-policy)#
```

Step 2 Specify the primary and secondary DNS servers:

```
hostname(config-group-policy)# dns-server value {ip_address [ip_address] | none}
hostname(config-group-policy)#
```

The first IP address specified is that of the primary DNS server. The second (optional) IP address is that of the secondary DNS server. Specifying the **none** keyword instead of an IP address sets DNS servers to a null value, which allows no DNS servers and prevents inheriting a value from a default or specified group policy.

Every time that you enter the **dns-server** command you overwrite the existing setting. For example, if you configure DNS server x.x.x.x and then configure DNS server y.y.y.y, the second command overwrites the first, and y.y.y.y becomes the sole DNS server. The same is true for multiple servers. To add a DNS server rather than overwrite previously configured servers, include the IP addresses of all DNS servers when you enter this command.

The following example shows how to configure DNS servers with the IP addresses 10.10.10.15, and 10.10.10.30 for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# dns-server value 10.10.10.15 10.10.10.30
hostname(config-group-policy)#
```

Step 3 Configure the DHCP network scope:

```
hostname(config-group-policy)# dhcp-network-scope {ip_address | none}
hostname(config-group-policy)#
```

DHCP scope specifies the range of IP addresses (that is, a subnetwork) that the security appliance DHCP server should use to assign addresses to users of this group policy.

The following example shows how to set an IP subnetwork of 10.10.85.0 (specifying the address range of 10.10.85.0 through 10.10.85.255) for the group policy named First Group:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# dhcp-network-scope 10.10.85.0
```

Configuring VPN-Specific Attributes

Follow the steps in this section to configure attributes that set the values of VPN attributes. These attributes control the access hours, the number of simultaneous logins allowed, the timeouts, the name of the ACL to use for VPN connections, and the tunnel protocol:

Step 1 Set the VPN access hours. To do this, you associate a group policy with a configured time-range policy, using the **vpn-access-hours** command in group-policy configuration mode.

```
hostname(config-group-policy)# vpn-access-hours value {time-range | none}
```

A group policy can inherit a time-range value from a default or specified group policy. To prevent this inheritance, enter the **none** keyword instead of the name of a time-range in this command. This keyword sets VPN access hours to a null value, which allows no time-range policy.

The time-range variable is the name of a set of access hours defined in global configuration mode using the **time-range** command. The following example shows how to associate the group policy named FirstGroup with a time-range policy called 824:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# vpn-access-hours value 824
```

- Step 2** Specify the number of simultaneous logins allowed for any user, using the **vpn-simultaneous-logins** command in group-policy configuration mode.

```
hostname(config-group-policy)# vpn-simultaneous-logins integer
```

The default value is 3. The range is an integer in the range 0 through 2147483647. A group policy can inherit this value from another group policy. Enter 0 to disable login and prevent user access. The following example shows how to allow a maximum of 4 simultaneous logins for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# vpn-simultaneous-logins 4
hostname(config-group-policy)#
```



Note While the maximum limit for the number of simultaneous logins is very large, allowing several could compromise security and affect performance.

- Step 3** Configure the user timeout period by entering the **vpn-idle-timeout** command in group-policy configuration mode or in username configuration mode:

```
hostname(config-group-policy)# vpn-idle-timeout {minutes | none}
hostname(config-group-policy)#
```

The minimum time is 1 minute, and the maximum time is 35791394 minutes. The default is 30 minutes. If there is no communication activity on the connection in this period, the security appliance terminates the connection.

A group policy can inherit this value from another group policy. To prevent inheriting a value, enter the **none** keyword instead of specifying a number of minutes with this command. The none keyword also permits an unlimited idle timeout period. It sets the idle timeout to a null value, thereby disallowing an idle timeout.

The following example shows how to set a VPN idle timeout of 15 minutes for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# vpn-idle-timeout 15
hostname(config-group-policy)#
```

- Step 4** Configure a maximum amount of time for VPN connections, using the **vpn-session-timeout** command in group-policy configuration mode or in username configuration mode.

```
hostname(config-group-policy)# vpn-session-timeout {minutes | none}
hostname(config-group-policy)#
```

The minimum time is 1 minute, and the maximum time is 35791394 minutes. There is no default value. At the end of this period of time, the security appliance terminates the connection.

A group policy can inherit this value from another group policy. To prevent inheriting a value, enter the **none** keyword instead of specifying a number of minutes with this command. Specifying the **none** keyword permits an unlimited session timeout period and sets session timeout with a null value, which disallows a session timeout.

The following example shows how to set a VPN session timeout of 180 minutes for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# vpn-session-timeout 180
hostname(config-group-policy)#
```

- Step 5** Specify the name of the ACL to use for VPN connections, using the **vpn-filter** command in group policy mode. (You can also configure this attribute in username mode, in which case the value configured under username supersedes the group-policy value.)

```
hostname(config-group-policy)# vpn-filter {value ACL name | none}
hostname(config-group-policy)#
```

You configure ACLs to permit or deny various types of traffic for this group policy. You then enter the **vpn-filter** command to apply those ACLs.

To remove the ACL, including a null value created by entering the **vpn-filter none** command, enter the **no** form of this command. The **no** option allows inheritance of a value from another group policy.

A group policy can inherit this value from another group policy. To prevent inheriting a value, enter the **none** keyword instead of specifying an ACL name. The **none** keyword indicates that there is no access list and sets a null value, thereby disallowing an access list.

The following example shows how to set a filter that invokes an access list named `acl_vpn` for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# vpn-filter acl_vpn
hostname(config-group-policy)#
```

- Step 6** Specify the VPN tunnel type (IPSec or WebVPN) for this group policy.

```
hostname(config-group-policy)# vpn-tunnel-protocol {webvpn | IPSec | l2tp-ipsec}
hostname(config-group-policy)#
```

The default is IPSec. To remove the attribute from the running configuration, enter the **no** form of this command.

```
hostname(config-group-policy)# no vpn-tunnel-protocol [webvpn | IPSec | l2tp-ipsec]
hostname(config-group-policy)#
```

The parameter values for this command follow:

- **IPSec**—Negotiates an IPSec tunnel between two peers (a remote access client or another secure gateway). Creates security associations that govern authentication, encryption, encapsulation, and key management.
- **webvpn**—Provides VPN services to remote users via an HTTPS-enabled web browser, and does not require a client.
- **l2tp-ipsec**—Negotiates an IPSec tunnel for an L2TP connection

Enter this command to configure one or more tunneling modes. You must configure at least one tunneling mode for users to connect over a VPN tunnel.

The following example shows how to configure the IPSec tunneling mode for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# vpn-tunnel-protocol IPSec
hostname(config-group-policy)#
```

Configuring Security Attributes

The attributes in this section specify certain security settings for the group:

- Step 1** Specify whether to let users store their login passwords on the client system, using the **password-storage** command with the **enable** keyword in group-policy configuration mode. To disable password storage, use the **password-storage** command with the **disable** keyword.

```
hostname(config-group-policy) # password-storage {enable | disable}
hostname(config-group-policy) #
```

For security reasons, password storage is disabled by default. Enable password storage only on systems that you know to be in secure sites.

To remove the password-storage attribute from the running configuration, enter the **no** form of this command:

```
hostname(config-group-policy) # no password-storage
hostname(config-group-policy) #
```

Specifying the **no** form enables inheritance of a value for password-storage from another group policy.

This command does not apply to interactive hardware client authentication or individual user authentication for hardware clients.

The following example shows how to enable password storage for the group policy named FirstGroup:

```
hostname(config) # group-policy FirstGroup attributes
hostname(config-group-policy) # password-storage enable
hostname(config-group-policy) #
```

- Step 2** Specify whether to enable IP compression, which is disabled by default.

```
hostname(config-group-policy) # ip-comp {enable | disable}
hostname(config-group-policy) #
```

To enable LZS IP compression, enter the **ip-comp** command with the **enable** keyword in group-policy configuration mode. To disable IP compression, enter the **ip-comp** command with the **disable** keyword.

To remove the **ip-comp** attribute from the running configuration, enter the **no** form of this command. This enables inheritance of a value from another group policy.

```
hostname(config-group-policy) # no ip-comp
hostname(config-group-policy) #
```

Enabling data compression might speed up data transmission rates for remote dial-in users connecting with modems.



Caution

Data compression increases the memory requirement and CPU usage for each user session and consequently decreases the overall throughput of the security appliance. For this reason, we recommend that you enable data compression only for remote users connecting with a modem. Design a group policy specific to modem users, and enable compression only for them.

- Step 3** Specify whether to require that users reauthenticate on IKE rekey by using the **re-xauth** command with the **enable** keyword in group-policy configuration mode. If you enable reauthentication on IKE rekey, the security appliance prompts the user to enter a username and password during initial Phase 1 IKE negotiation and also prompts for user authentication whenever an IKE rekey occurs. Reauthentication provides additional security.

If the configured rekey interval is very short, users might find the repeated authorization requests inconvenient. To avoid repeated authorization requests, disable reauthentication. To check the configured rekey interval, in monitoring mode, enter the **show crypto ipsec sa** command to view the security association lifetime in seconds and lifetime in kilobytes of data. To disable user reauthentication on IKE rekey, enter the **disable** keyword. Reauthentication on IKE rekey is disabled by default.

```
hostname(config-group-policy)# re-xauth {enable | disable}
hostname(config-group-policy)#
```

To enable inheritance of a value for reauthentication on IKE rekey from another group policy, remove the re-xauth attribute from the running configuration by entering the **no** form of this command.

```
hostname(config-group-policy)# no re-xauth
hostname(config-group-policy)#
```



Note Reauthentication fails if there is no user at the other end of the connection.

- Step 4** Specify whether to restrict remote users to access only through the tunnel group, using the **group-lock** command in group-policy configuration mode.

```
hostname(config-group-policy)# group-lock {value tunnel-grp-name | none}
hostname(config-group-policy)# no group-lock
hostname(config-group-policy)#
```

The *tunnel-grp-name* variable specifies the name of an existing tunnel group that the security appliance requires for the user to connect. Group-lock restricts users by checking if the group configured in the VPN client is the same as the tunnel group to which the user is assigned. If it is not, the security appliance prevents the user from connecting. If you do not configure group-lock, the security appliance authenticates users without regard to the assigned group. Group locking is disabled by default.

To remove the **group-lock** attribute from the running configuration, enter the **no** form of this command. This option allows inheritance of a value from another group policy.

To disable group-lock, enter the **group-lock** command with the **none** keyword. The none keyword sets group-lock to a null value, thereby allowing no group-lock restriction. It also prevents inheriting a group-lock value from a default or specified group policy.

- Step 5** Specify whether to enable perfect forward secrecy. In IPSec negotiations, perfect forward secrecy ensures that each new cryptographic key is unrelated to any previous key. A group policy can inherit a value for perfect forward secrecy from another group policy. Perfect forward secrecy is disabled by default. To enable perfect forward secrecy, use the **pfs** command with the **enable** keyword in group-policy configuration mode.

```
hostname(config-group-policy)# pfs {enable | disable}
hostname(config-group-policy)#
```

To disable perfect forward secrecy, enter the **pfs** command with the **disable** keyword.

To remove the perfect forward secrecy attribute from the running configuration and prevent inheriting a value, enter the **no** form of this command.

```
hostname(config-group-policy)# no pfs
hostname(config-group-policy)#
```

Configuring the Banner Message

Specify the banner, or welcome message, if any, that you want to display. The default is no banner. The message that you specify is displayed on remote clients when they connect. To specify a banner, enter the **banner** command in group-policy configuration mode. The banner text can be up to 510 characters long. Enter the “\n” sequence to insert a carriage return.

**Note**

A carriage-return/line-feed included in the banner counts as two characters.

To delete a banner, enter the **no** form of this command. Be aware that using the **no** version of the command deletes all banners for the group policy.

A group policy can inherit this value from another group policy. To prevent inheriting a value, enter the **none** keyword instead of specifying a value for the banner string, as follows:

```
hostname(config-group-policy)# banner {value banner_string | none}
```

The following example shows how to create a banner for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes  
hostname(config-group-policy)# banner value Welcome to Cisco Systems 7.0.
```

Configuring IPSec-UDP Attributes

IPSec over UDP, sometimes called IPSec through NAT, lets a Cisco VPN client or hardware client connect via UDP to a security appliance that is running NAT. It is disabled by default. IPSec over UDP is proprietary; it applies only to remote-access connections, and it requires mode configuration. The security appliance exchanges configuration parameters with the client while negotiating SAs. Using IPSec over UDP may slightly degrade system performance.

To enable IPSec over UDP, configure the **ipsec-udp** command with the **enable** keyword in group-policy configuration mode, as follows:

```
hostname(config-group-policy)# ipsec-udp {enable | disable}  
hostname(config-group-policy)# no ipsec-udp
```

To use IPSec over UDP, you must also configure the **ipsec-udp-port** command, as described below.

To disable IPSec over UDP, enter the **disable** keyword. To remove the IPSec over UDP attribute from the running configuration, enter the **no** form of this command. This enables inheritance of a value for IPSec over UDP from another group policy.

The Cisco VPN client must also be configured to use IPSec over UDP (it is configured to use it by default). The VPN 3002 requires no configuration to use IPSec over UDP.

The following example shows how to set IPSec over UDP for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes  
hostname(config-group-policy)# ipsec-udp enable
```

If you enabled IPSec over UDP, you must also configure the **ipsec-udp-port** command in group-policy configuration mode. This command sets a UDP port number for IPSec over UDP. In IPSec negotiations, the security appliance listens on the configured port and forwards UDP traffic for that port even if other filter rules drop UDP traffic. The port numbers can range from 4001 through 49151. The default port value is 10000.

To disable the UDP port, enter the **no** form of this command. This enables inheritance of a value for the IPSec over UDP port from another group policy.

```
hostname(config-group-policy)# ipsec-udp-port port
```

The following example shows how to set an IPSec UDP port to port 4025 for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# ipsec-udp-port 4025
```

Configuring Split-Tunneling Attributes

Split tunneling lets a remote-access IPSec client conditionally direct packets over an IPSec tunnel in encrypted form or to a network interface in clear text form. With split tunneling enabled, packets not bound for destinations on the other side of the IPSec tunnel do not have to be encrypted, sent across the tunnel, decrypted, and then routed to a final destination. This command applies this split tunneling policy to a specific network.

Setting the Split-Tunneling Policy

Set the rules for tunneling traffic by specifying the split-tunneling policy:

```
hostname(config-group-policy)# split-tunnel-policy {tunnelall | tunnelspecified |
excludespecified}
hostname(config-group-policy)# no split-tunnel-policy
```

The default is to tunnel all traffic. To set a split tunneling policy, enter the **split-tunnel-policy** command in group-policy configuration mode. To remove the **split-tunnel-policy** attribute from the running configuration, enter the **no** form of this command. This enables inheritance of a value for split tunneling from another group policy.

The **excludespecified** keyword defines a list of networks to which traffic goes in the clear. This feature is useful for remote users who want to access devices on their local network, such as printers, while they are connected to the corporate network through a tunnel. This option applies only to the Cisco VPN client.

The **tunnelall** keyword specifies that no traffic goes in the clear or to any other destination than the security appliance. This, in effect, disables split tunneling. Remote users reach Internet networks through the corporate network and do not have access to local networks. This is the default option.

The **tunnelspecified** keyword tunnels all traffic from or to the specified networks. This option enables split tunneling. It lets you create a network list of addresses to tunnel. Data to all other addresses travels in the clear and is routed by the remote user's Internet service provider.



Note

Split tunneling is primarily a traffic management feature, not a security feature. For optimum security, we recommend that you do not enable split tunneling.

The following example shows how to set a split tunneling policy of tunneling only specified networks for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# split-tunnel-policy tunnelspecified
```

Creating a Network List for Split-Tunneling

Create a network list for split tunneling using the **split-tunnel-network-list** command in group-policy configuration mode.

```
hostname(config-group-policy)# split-tunnel-network-list {value access-list_name | none}
hostname(config-group-policy)# no split-tunnel-network-list value [access-list_name]
```

Split tunneling network lists distinguish networks that require traffic to travel across the tunnel from those that do not require tunneling. The security appliance makes split tunneling decisions on the basis of a network list, which is an ACL that consists of a list of addresses on the private network. Only standard-type ACLs are allowed.

The **value** *access-list name* parameter identifies an access list that enumerates the networks to tunnel or not tunnel.

The **none** keyword indicates that there is no network list for split tunneling; the security appliance tunnels all traffic. Specifying the **none** keyword sets a split tunneling network list with a null value, thereby disallowing split tunneling. It also prevents inheriting a default split tunneling network list from a default or specified group policy.

To delete a network list, enter the **no** form of this command. To delete all split tunneling network lists, enter the **no split-tunnel-network-list** command without arguments. This command deletes all configured network lists, including a null list if you created one by entering the **none** keyword.

When there are no split tunneling network lists, users inherit any network lists that exist in the default or specified group policy. To prevent users from inheriting such network lists, enter the **split-tunnel-network-list none** command.

The following example shows how to set a network list called FirstList for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# split-tunnel-network-list FirstList
```

Configuring Domain Attributes for Tunneling

You can specify a default domain name for tunneled packets or a list of domains to be resolved through the split tunnel. The following sections describe how to set these domains.

Defining a Default Domain Name for Tunneled Packets

The security appliance passes the default domain name to the IPsec client to append to DNS queries that omit the domain field. When there are no default domain names, users inherit the default domain name in the default group policy. To specify the default domain name for users of the group policy, enter the **default-domain** command in group-policy configuration mode. To delete a domain name, enter the **no** form of this command.

```
hostname(config-group-policy)# default-domain {value domain-name | none}
hostname(config-group-policy)# no default-domain [domain-name]
```

The **value** *domain-name* parameter identifies the default domain name for the group. To specify that there is no default domain name, enter the **none** keyword. This command sets a default domain name with a null value, which disallows a default domain name and prevents inheriting a default domain name from a default or specified group policy.

To delete all default domain names, enter the **no default-domain** command without arguments. This command deletes all configured default domain names, including a null list if you created one by entering the **default-domain** command with the **none** keyword. The **no** form allows inheriting a domain name.

The following example shows how to set a default domain name of FirstDomain for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# default-domain value FirstDomain
```

Defining a List of Domains for Split Tunneling

Enter a list of domains to be resolved through the split tunnel. Enter the **split-dns** command in group-policy configuration mode. To delete a list, enter the **no** form of this command.



Note

The AnyConnect VPN client and the SSL VPN Client do not support split DNS.

When there are no split tunneling domain lists, users inherit any that exist in the default group policy. To prevent users from inheriting such split tunneling domain lists, enter the **split-dns** command with the **none** keyword.

To delete all split tunneling domain lists, enter the **no split-dns** command without arguments. This deletes all configured split tunneling domain lists, including a null list created by issuing the **split-dns** command with the **none** keyword.

The parameter **value domain-name** provides a domain name that the security appliance resolves through the split tunnel. The **none** keyword indicates that there is no split DNS list. It also sets a split DNS list with a null value, thereby disallowing a split DNS list, and prevents inheriting a split DNS list from a default or specified group policy. The syntax of the command is as follows:

```
hostname(config-group-policy)# split-dns {value domain-name1 [domain-name2...
domain-nameN] | none}
hostname(config-group-policy)# no split-dns [domain-name domain-name2 domain-nameN]
```

Enter a single space to separate each entry in the list of domains. There is no limit on the number of entries, but the entire string can be no longer than 255 characters. You can use only alphanumeric characters, hyphens (-), and periods (.). If the default domain name is to be resolved through the tunnel, you must explicitly include that name in this list.

The following example shows how to configure the domains Domain1, Domain2, Domain3, and Domain4 to be resolved through split tunneling for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# split-dns value Domain1 Domain2 Domain3 Domain4
```

Configuring DHCP Intercept

A Microsoft XP anomaly results in the corruption of domain names if split tunnel options exceed 255 bytes. To avoid this problem, the security appliance limits the number of routes it sends to 27 to 40 routes, with the number of routes dependent on the classes of the routes.

DHCP Intercept lets Microsoft Windows XP clients use split-tunneling with the security appliance. The security appliance replies directly to the Microsoft Windows XP client DHCP Inform message, providing that client with the subnet mask, domain name, and classless static routes for the tunnel IP address. For Windows clients prior to Windows XP, DHCP Intercept provides the domain name and subnet mask. This is useful in environments in which using a DHCP server is not advantageous.

The **intercept-dhcp** command enables or disables DHCP intercept. The syntax of this command is as follows:

[no] intercept-dhcp

```
hostname(config-group-policy)# intercept-dhcp netmask {enable | disable}
```



```
hostname(config-group-policy) #
```

The *netmask* variable provides the subnet mask for the tunnel IP address. The **no** version of the command removes the DHCP intercept from the configuration.

The following example shows how to set DHCP Intercepts for the group policy named FirstGroup:

```
hostname(config) # group-policy FirstGroup attributes  
hostname(config-group-policy) # intercept-dhcp enable
```

Configuring Attributes for VPN Hardware Clients

The commands in this section enable or disable secure unit authentication and user authentication, and set a user authentication timeout value for VPN hardware clients. They also let you allow Cisco IP phones and LEAP packets to bypass individual user authentication and allow hardware clients using Network Extension Mode to connect.

Configuring Secure Unit Authentication

Secure unit authentication provides additional security by requiring VPN hardware clients to authenticate with a username and password each time that the client initiates a tunnel. With this feature enabled, the hardware client does not have a saved username and password. Secure unit authentication is disabled by default.



Note

With this feature enabled, to bring up a VPN tunnel, a user must be present to enter the username and password.

Secure unit authentication requires that you have an authentication server group configured for the tunnel group the hardware client(s) use. If you require secure unit authentication on the primary security appliance, be sure to configure it on any backup servers as well.

Specify whether to enable secure unit authentication by entering the **secure-unit-authentication** command with the **enable** keyword in group-policy configuration mode.

```
hostname(config-group-policy) # secure-unit-authentication {enable | disable}  
hostname(config-group-policy) # no secure-unit-authentication
```

To disable secure unit authentication, enter the **disable** keyword. To remove the secure unit authentication attribute from the running configuration, enter the **no** form of this command. This option allows inheritance of a value for secure unit authentication from another group policy.

The following example shows how to enable secure unit authentication for the group policy named FirstGroup:

```
hostname(config) # group-policy FirstGroup attributes  
hostname(config-group-policy) # secure-unit-authentication enable
```

Configuring User Authentication

User authentication is disabled by default. When enabled, user authentication requires that individual users behind a hardware client authenticate to gain access to the network across the tunnel. Individual users authenticate according to the order of authentication servers that you configure.

Specify whether to enable user authentication by entering the **user-authentication** command with the **enable** keyword in group-policy configuration mode.

```
hostname(config-group-policy)# user-authentication {enable | disable}
hostname(config-group-policy)# no user-authentication
```

To disable user authentication, enter the **disable** keyword. To remove the user authentication attribute from the running configuration, enter the **no** form of this command. This option allows inheritance of a value for user authentication from another group policy.

If you require user authentication on the primary security appliance, be sure to configure it on any backup servers as well.

The following example shows how to enable user authentication for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# user-authentication enable
```

Configuring an Idle Timeout

Set an idle timeout for individual users behind hardware clients by entering the **user-authentication-idle-timeout** command in group-policy configuration mode. If there is no communication activity by a user behind a hardware client in the idle timeout period, the security appliance terminates the client's access:

```
hostname(config-group-policy)# user-authentication-idle-timeout {minutes | none}
hostname(config-group-policy)# no user-authentication-idle-timeout
```



Note

The **user-authentication-idle-timeout** command terminates only the client's access through the VPN tunnel, not the VPN tunnel itself.

The *minutes* parameter specifies the number of minutes in the idle timeout period. The minimum is 1 minute, the default is 30 minutes, and the maximum is 35791394 minutes.

To delete the idle timeout value, enter the **no** form of this command. This option allows inheritance of an idle timeout value from another group policy.

To prevent inheriting an idle timeout value, enter the **user-authentication-idle-timeout** command with the **none** keyword. This command sets the idle timeout with a null value, which disallows an idle timeout and prevents inheriting an user authentication idle timeout value from a default or specified group policy.

The following example shows how to set an idle timeout value of 45 minutes for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# user-authentication-idle-timeout 45
```

Configuring IP Phone Bypass

You can allow Cisco IP phones to bypass individual user authentication behind a hardware client. To enable IP Phone Bypass, enter the **ip-phone-bypass** command with the **enable** keyword in group-policy configuration mode. IP Phone Bypass lets IP phones behind hardware clients connect without undergoing user authentication processes. IP Phone Bypass is disabled by default. If enabled, secure unit authentication remains in effect.

To disable IP Phone Bypass, enter the **disable** keyword. To remove the IP phone Bypass attribute from the running configuration, enter the **no** form of this command. This option allows inheritance of a value for IP Phone Bypass from another group policy:

```
hostname(config-group-policy) # ip-phone-bypass {enable | disable}  
hostname(config-group-policy) # no ip-phone-bypass
```

Configuring LEAP Bypass

When LEAP Bypass is enabled, LEAP packets from wireless devices behind a VPN 3002 hardware client travel across a VPN tunnel prior to user authentication. This action lets workstations using Cisco wireless access point devices establish LEAP authentication and then authenticate again per user authentication. LEAP Bypass is disabled by default.

To allow LEAP packets from Cisco wireless access points to bypass individual users authentication, enter the **leap-bypass** command with the **enable** keyword in group-policy configuration mode. To disable LEAP Bypass, enter the **disable** keyword. To remove the LEAP Bypass attribute from the running configuration, enter the **no** form of this command. This option allows inheritance of a value for LEAP Bypass from another group policy:

```
hostname(config-group-policy) # leap-bypass {enable | disable}  
hostname(config-group-policy) # no leap-bypass
```



Note

IEEE 802.1X is a standard for authentication on wired and wireless networks. It provides wireless LANs with strong mutual authentication between clients and authentication servers, which can provide dynamic per-user, per session wireless encryption privacy (WEP) keys, removing administrative burdens and security issues that are present with static WEP keys.

Cisco Systems has developed an 802.1X wireless authentication type called Cisco LEAP. LEAP (Lightweight Extensible Authentication Protocol) implements mutual authentication between a wireless client on one side of a connection and a RADIUS server on the other side. The credentials used for authentication, including a password, are always encrypted before they are transmitted over the wireless medium.

Cisco LEAP authenticates wireless clients to RADIUS servers. It does not include RADIUS accounting services.

This feature does not work as intended if you enable interactive hardware client authentication.



Caution

There might be security risks to your network in allowing any unauthenticated traffic to traverse the tunnel.

The following example shows how to set LEAP Bypass for the group policy named FirstGroup:

```
hostname(config) # group-policy FirstGroup attributes  
hostname(config-group-policy) # leap-bypass enable
```

Enabling Network Extension Mode

Network extension mode lets hardware clients present a single, routable network to the remote private network over the VPN tunnel. IPSec encapsulates all traffic from the private network behind the hardware client to networks behind the security appliance. PAT does not apply. Therefore, devices behind the security appliance have direct access to devices on the private network behind the hardware client over the tunnel, and only over the tunnel, and vice versa. The hardware client must initiate the tunnel, but after the tunnel is up, either side can initiate data exchange.

Enable network extension mode for hardware clients by entering the **nem** command with the **enable** keyword in group-policy configuration mode:

```
hostname(config-group-policy)# nem {enable | disable}
hostname(config-group-policy)# no nem
```

To disable NEM, enter the **disable** keyword. To remove the NEM attribute from the running configuration, enter the **no** form of this command. This option allows inheritance of a value from another group policy.

The following example shows how to set NEM for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# nem enable
```

Configuring Backup Server Attributes

Configure backup servers if you plan on using them. IPsec backup servers let a VPN client connect to the central site when the primary security appliance is unavailable. When you configure backup servers, the security appliance pushes the server list to the client as the IPsec tunnel is established. Backup servers do not exist until you configure them, either on the client or on the primary security appliance.

Configure backup servers either on the client or on the primary security appliance. If you configure backup servers on the security appliance, it pushes the backup server policy to the clients in the group, replacing the backup server list on the client if one is configured.



Note

If you are using hostnames, it is wise to have backup DNS and WINS servers on a separate network from that of the primary DNS and WINS servers. Otherwise, if clients behind a hardware client obtain DNS and WINS information from the hardware client via DHCP, and the connection to the primary server is lost, and the backup servers have different DNS and WINS information, clients cannot be updated until the DHCP lease expires. In addition, if you use hostnames and the DNS server is unavailable, significant delays can occur.

To configure backup servers, enter the **backup-servers** command in group-policy configuration mode:

```
hostname(config-group-policy)# backup-servers {server1 server2... server10 |  
clear-client-config | keep-client-config}
```

To remove a backup server, enter the **no** form of this command with the backup server specified. To remove the backup-servers attribute from the running configuration and enable inheritance of a value for backup-servers from another group policy, enter the **no** form of this command without arguments.

```
hostname(config-group-policy)# no backup-servers [server1 server2... server10 |  
clear-client-config | keep-client-config]
```

The **clear-client-config** keyword specifies that the client uses no backup servers. The security appliance pushes a null server list.

The **keep-client-config** keyword specifies that the security appliance sends no backup server information to the client. The client uses its own backup server list, if configured. This is the default.

The *server1 server 2.... server10* parameter list is a space-delimited, priority-ordered list of servers for the VPN client to use when the primary security appliance is unavailable. This list identifies servers by IP address or hostname. The list can be 500 characters long, and it can contain up to 10 entries.

The following example shows how to configure backup servers with IP addresses 10.10.10.1 and 192.168.10.14, for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# backup-servers 10.10.10.1 192.168.10.14
```

Configuring Microsoft Internet Explorer Client Parameters

The following commands configure the proxy server parameters for a Microsoft Internet Explorer client.

- Step 1** Configure a Microsoft Internet Explorer browser proxy server and port for a client PC by entering the **msie-proxy server** command in group-policy configuration mode:

```
hostname(config-group-policy)# msie-proxy server {value server[:port] | none}
hostname(config-group-policy)#
```

The default value is **none**. To remove the attribute from the configuration, use the **no** form of the command.

```
hostname(config-group-policy)# no msie-proxy server
hostname(config-group-policy)#
```

The line containing the proxy server IP address or hostname and the port number must be less than 100 characters long.

The following example shows how to configure the IP address 192.168.10.1 as a Microsoft Internet Explorer proxy server, using port 880, for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# msie-proxy server value 192.168.21.1:880
hostname(config-group-policy)#
```

- Step 2** Configure the Microsoft Internet Explorer browser proxy actions (“methods”) for a client PC by entering the **msie-proxy method** command in group-policy configuration mode.

```
hostname(config-group-policy)# msie-proxy method [auto-detect | no-modify | no-proxy |
use-server]
hostname(config-group-policy)#
```

The default value is **use-server**. To remove the attribute from the configuration, use the **no** form of the command.

```
hostname(config-group-policy)# no msie-proxy method [auto-detect | no-modify | no-proxy |
use-server]
hostname(config-group-policy)#
```

The available methods are as follows:

- **auto-detect**—Enables the use of automatic proxy server detection in Internet Explorer for the client PC.
- **no-modify**—Leaves the HTTP browser proxy server setting in Internet Explorer unchanged for this client PC.
- **no-proxy**—Disables the HTTP proxy setting in Internet Explorer for the client PC.
- **use-server**—Sets the HTTP proxy server setting in Internet Explorer to use the value configured in the **msie-proxy server** command.

The line containing the proxy server IP address or hostname and the port number must be less than 100 characters long.

The following example shows how to configure auto-detect as the Microsoft Internet Explorer proxy setting for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# msie-proxy method auto-detect
hostname(config-group-policy)#
```

The following example configures the Microsoft Internet Explorer proxy setting for the group policy named FirstGroup to use the server QAsrver, port 1001 as the server for the client PC:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# msie-proxy server QAsrver:port 1001
hostname(config-group-policy)# msie-proxy method use-server
hostname(config-group-policy)#
```

Step 3 Configure Microsoft Internet Explorer browser proxy exception list settings for a local bypass on the client PC by entering the **msie-proxy except-list** command in group-policy configuration mode. These addresses are not accessed by a proxy server. This list corresponds to the Exceptions box in the Proxy Settings dialog box in Internet Explorer.

```
hostname(config-group-policy)# msie-proxy except-list {value server[:port] | none}
hostname(config-group-policy)#
```

To remove the attribute from the configuration, use the **no** form of the command.

```
hostname(config-group-policy)# no msie-proxy except-list
hostname(config-group-policy)#
```

- **value server:port**—Specifies the IP address or name of an MSIE server and port that is applied for this client PC. The port number is optional.
- **none**—Indicates that there is no IP address/hostname or port and prevents inheriting an exception list.

By default, msie-proxy except-list is disabled.

The line containing the proxy server IP address or hostname and the port number must be less than 100 characters long.

The following example shows how to set a Microsoft Internet Explorer proxy exception list, consisting of the server at IP address 192.168.20.1, using port 880, for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# msie-proxy except-list value 192.168.20.1:880
hostname(config-group-policy)#
```

Step 4 Enable or disable Microsoft Internet Explorer browser proxy local-bypass settings for a client PC by entering the **msie-proxy local-bypass** command in group-policy configuration mode.

```
hostname(config-group-policy)# msie-proxy local-bypass {enable | disable}
hostname(config-group-policy)#
```

To remove the attribute from the configuration, use the **no** form of the command.

```
hostname(config-group-policy)# no msie-proxy local-bypass {enable | disable}
hostname(config-group-policy)#
```

By default, msie-proxy local-bypass is disabled.

The following example shows how to enable Microsoft Internet Explorer proxy local-bypass for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# msie-proxy local-bypass enable
hostname(config-group-policy)#
```

Configuring Network Admission Control Parameters

The group-policy NAC commands in this section all have default values. Unless you have a good reason for changing them, accept the default values for these parameters.

The security appliance uses Extensible Authentication Protocol (EAP) over UDP (EAPoUDP) messaging to validate the posture of remote hosts. Posture validation involves the checking of a remote host for compliancy with safety requirements before the assignment of a network access policy. An Access Control Server must be configured for Network Admission Control before you configure NAC on the security appliance.

The Access Control Server downloads the posture token, an informational text string configurable on the ACS, to the security appliance to aid in system monitoring, reporting, debugging, and logging. A typical posture token is Healthy, Checkup, Quarantine, Infected, or Unknown. Following posture validation or clientless authentication, the ACS downloads the access policy for the session to the security appliance.

The following parameters let you configure Network Admission Control settings for the default group policy or an alternative group policy.

Step 1 *(Optional)* Configure the status query timer period. The security appliance starts the status query timer after each successful posture validation and status query response. The expiration of this timer triggers a query for changes in the host posture, referred to as a status query. Enter the number of seconds in the range 30 through 1800. The default setting is 300.

To specify the interval between each successful posture validation in a Network Admission Control session and the next query for changes in the host posture, use the **nac-sq-period** command in group-policy configuration mode:

```
hostname(config-group-policy)# nac-sq-period seconds
hostname(config-group-policy)#
```

To inherit the value of the status query timer from the default group policy, access the alternative group policy from which to inherit it, then use the **no** form of this command:

```
hostname(config-group-policy)# no nac-sq-period [seconds]
hostname(config-group-policy)#
```

The following example changes the value of the status query timer to 1800 seconds:

```
hostname(config-group-policy)# nac-sq-period 1800
hostname(config-group-policy)
```

The following example inherits the value of the status query timer from the default group policy:

```
hostname(config-group-policy)# no nac-sq-period
hostname(config-group-policy)#
```

Step 2 *(Optional)* Configure the NAC revalidation period. The security appliance starts the revalidation timer after each successful posture validation. The expiration of this timer triggers the next unconditional posture validation. The security appliance maintains posture validation during revalidation. The default group policy becomes effective if the Access Control Server is unavailable during posture validation or revalidation. Enter the interval in seconds between each successful posture validation. The range is 300 through 86400. The default setting is 36000.

To specify the interval between each successful posture validation in a Network Admission Control session, use the **nac-reval-period** command in group-policy configuration mode:

```
hostname(config-group-policy)# nac-reval-period seconds
hostname(config-group-policy)#
```

To inherit the value of the Revalidation Timer from the default group policy, access the alternative group policy from which to inherit it, then use the **no** form of this command:

```
hostname(config-group-policy) # no nac-reval-period [seconds]
hostname(config-group-policy) #
```

The following example changes the revalidation timer to 86400 seconds:

```
hostname(config-group-policy) # nac-reval-period 86400
hostname(config-group-policy) #
```

The following example inherits the value of the revalidation timer from the default group policy:

```
hostname(config-group-policy) # no nac-reval-period
hostname(config-group-policy) #
```

- Step 3** (Optional) Configure the default ACL for NAC. The security appliance applies the security policy associated with the selected ACL if posture validation fails. Specify **none** or an extended ACL. The default setting is **none**. If the setting is **none** and posture validation fails, the security appliance applies the default group policy.

To specify the ACL to be used as the default ACL for Network Admission Control sessions that fail posture validation, use the **nac-default-acl** command in group-policy configuration mode:

```
hostname(config-group-policy) # nac-default-acl {acl-name | none}
hostname(config-group-policy) #
```

To inherit the ACL from the default group policy, access the alternative group policy from which to inherit it, then use the **no** form of this command:

```
hostname(config-group-policy) # no nac-default-acl [acl-name | none]
hostname(config-group-policy) #
```

The elements of this command are as follows:

- **acl-name**—Specifies the name of the posture validation server group, as configured on the security appliance using the **aaa-server host** command. The name must match the server-tag variable specified in that command.
- **none**—Disables inheritance of the ACL from the default group policy and does not apply an ACL to NAC sessions that fail posture validation.

Because NAC is disabled by default, VPN traffic traversing the security appliance is not subject to the NAC Default ACL until NAC is enabled.

The following example identifies acl-1 as the ACL to be applied when posture validation fails:

```
hostname(config-group-policy) # nac-default-acl acl-1
hostname(config-group-policy) #
```

The following example inherits the ACL from the default group policy:

```
hostname(config-group-policy) # no nac-default-acl
hostname(config-group-policy) #
```

The following example disables inheritance of the ACL from the default group policy and does not apply an ACL to NAC sessions that fail posture validation:

```
hostname(config-group-policy) # nac-default-acl none
hostname(config-group-policy) #
```

- Step 4** Configure NAC exemptions for VPN. By default, the exemption list is empty. The default value of the filter attribute is **none**. Enter the **vpn-nac-exempt** once for each operating system (and ACL) to be matched to exempt remote hosts from posture validation.

To add an entry to the list of remote computer types that are exempt from posture validation, use the **vpn-nac-exempt** command in group-policy configuration mode.

```
hostname(config-group-policy)# vpn-nac-exempt os "os name" [filter {acl-name | none}]
[disable]
hostname(config-group-policy)#
```

To disable inheritance and specify that all hosts are subject to posture validation, use the **none** keyword immediately following **vpn-nac-exempt**.

```
hostname(config-group-policy)# vpn-nac-exempt none
hostname(config-group-policy)#
```

To remove an entry from the exemption list, use the **no** form of this command and name the operating system (and ACL) in the entry to be removed.

```
hostname(config-group-policy)# no vpn-nac-exempt [os "os name"] [filter {acl-name | none}]
[disable]
hostname(config-group-policy)#
```

To remove all entries from the exemption list associated with this group policy and inherit the list from the default group policy, use the **no** form of this command without specifying additional keywords.

```
hostname(config-group-policy)# no vpn-nac-exempt
hostname(config-group-policy)#
```

The syntax elements for these commands are as follows:

- **acl-name**—Name of the ACL present in the security appliance configuration.
- **disable**—Disables the entry in the exemption list without removing it from the list.
- **filter**—(*Optional*) filter to apply an ACL to filter the traffic if the computer matches the OS name.
- **none**—When entered immediately after **vpn-nac-exempt**, this keyword disables inheritance and specifies that all hosts will be subject to posture validation. When entered immediately after **filter**, this keyword indicates that the entry does not specify an ACL.
- **OS**—Exempts an operating system from posture validation.
- **os name**—Operating system name. Quotation marks are required only if the name includes a space (for example, "Windows XP").

The following example adds all hosts running Windows XP to the list of computers that are exempt from posture validation:

```
hostname(config-group-policy)# vpn-nac-exempt os "Windows XP"
hostname(config-group-policy)
```

The following example exempts all hosts running Windows 98 that match an ACE in the ACL named **acl-1**:

```
hostname(config-group-policy)# vpn-nac-exempt os "Windows 98" filter acl-1
hostname(config-group-policy)
```

The following example adds the same entry to the exemption list, but disables it:

```
hostname(config-group-policy)# vpn-nac-exempt os "Windows 98" filter acl-1 disable
hostname(config-group-policy)
```

The following example removes the same entry from the exemption list, regardless of whether it is disabled:

```
hostname(config-group-policy)# no vpn-nac-exempt os "Windows 98" filter acl-1
hostname(config-group-policy)
```

The following example disables inheritance and specifies that all hosts will be subject to posture validation:

```
hostname(config-group-policy)# no vpn-nac-exempt none
hostname(config-group-policy)
```

The following example removes all entries from the exemption list:

```
hostname(config-group-policy)# no vpn-nac-exempt
hostname(config-group-policy)
```

Step 5 Enable or disable Network Admission Control by entering the following command:

```
hostname(config-group-policy)# nac {enable | disable}
hostname(config-group-policy)#
```

To inherit the NAC setting from the default group policy, access the alternative group policy from which to inherit it, then use the **no** form of this command:

```
hostname(config-group-policy)# no nac [enable | disable]
hostname(config-group-policy)#
```

By default, NAC is disabled. Enabling NAC requires posture validation for remote access. If the remote computer passes the validation checks, the ACS server downloads the access policy for the security appliance to enforce. NAC is disabled by default.

An Access Control Server must be present on the network.

The following example enables NAC for the group policy:

```
hostname(config-group-policy)# nac enable
hostname(config-group-policy)#
```

Configuring Address Pools

Configure a list of address pools for allocating addresses to remote clients by entering the **address-pools** command in group-policy attributes configuration mode:

```
hostname(config-group-policy)# address-pools value address_pool1 [...address_pool6]
hostname(config-group-policy)#
```

The address-pools settings in this command override the local pool settings in the group. You can specify a list of up to six local address pools to use for local address allocation.

The order in which you specify the pools is significant. The security appliance allocates addresses from these pools in the order in which the pools appear in this command.

To remove the attribute from the group policy and enable inheritance from other sources of group policy, use the **no** form of this command:

```
hostname(config-group-policy)# no address-pools value address_pool1 [...address_pool6]
hostname(config-group-policy)#
```

The command **address-pools none** disables this attribute from being inherited from other sources of policy, such as the DefaultGrpPolicy:

```
hostname(config-group-policy)# address-pools none
hostname(config-group-policy)#
```

The command **no address pools none** removes the **address-pools none** command from the configuration, restoring the default value, which is to allow inheritance.

```
hostname(config-group-policy)# no address-pools none
hostname(config-group-policy)#
```

The syntax elements of this command are as follows:

- **address_pool**—Specifies the name of the address pool configured with the **ip local pool** command. You can specify up to 6 local address pools.
- **none**—Specifies that no address pools are configured and disables inheritance from other sources of group policy.
- **value**—Specifies a list of up to 6 address pools from which to assign addresses.

The following example entered in config-general configuration mode, configures pool 1 and pool20 as lists of address pools to use for allocating addresses to remote clients for GroupPolicy1:

```
hostname(config)# ip local pool pool1 192.168.10.1-192.168.10.100 mask 255.255.0.0
hostname(config)# ip local pool pool20 192.168.20.1-192.168.20.200 mask 255.255.0.0
hostname(config)# group-policy GroupPolicy1 attributes
hostname(config-group-policy)# address-pools value pool1 pool20
hostname(config-group-policy)#
```

Configuring Firewall Policies

A *firewall* isolates and protects a computer from the Internet by inspecting each inbound and outbound individual packet of data to determine whether to allow or drop it. Firewalls provide extra security if remote users in a group have split tunneling configured. In this case, the firewall protects the user's PC, and thereby the corporate network, from intrusions by way of the Internet or the user's local LAN. Remote users connecting to the security appliance with the VPN client can choose the appropriate firewall option.

Set personal firewall policies that the security appliance pushes to the VPN client during IKE tunnel negotiation by using the **client-firewall** command in group-policy configuration mode. To delete a firewall policy, enter the **no** form of this command.

To delete all firewall policies, enter the **no client-firewall** command without arguments. This command deletes all configured firewall policies, including a null policy if you created one by entering the **client-firewall** command with the **none** keyword.

When there are no firewall policies, users inherit any that exist in the default or other group policy. To prevent users from inheriting such firewall policies, enter the **client-firewall** command with the **none** keyword.

The Add or Edit Group Policy window, Client Firewall tab, lets you configure firewall settings for VPN clients for the group policy being added or modified.



Note

Only VPN clients running Microsoft Windows can use these firewall features. They are currently not available to hardware clients or other (non-Windows) software clients.

In the first scenario, a remote user has a personal firewall installed on the PC. The VPN client enforces firewall policy defined on the local firewall, and it monitors that firewall to make sure it is running. If the firewall stops running, the VPN client drops the connection to the security appliance. (This firewall enforcement mechanism is called *Are You There (AYT)*, because the VPN client monitors the firewall by

sending it periodic “are you there?” messages; if no reply comes, the VPN client knows the firewall is down and terminates its connection to the security appliance.) The network administrator might configure these PC firewalls originally, but with this approach, each user can customize his or her own configuration.

In the second scenario, you might prefer to enforce a centralized firewall policy for personal firewalls on VPN client PCs. A common example would be to block Internet traffic to remote PCs in a group using split tunneling. This approach protects the PCs, and therefore the central site, from intrusions from the Internet while tunnels are established. This firewall scenario is called *push policy* or *Central Protection Policy (CPP)*. On the security appliance, you create a set of traffic management rules to enforce on the VPN client, associate those rules with a filter, and designate that filter as the firewall policy. The security appliance pushes this policy down to the VPN client. The VPN client then in turn passes the policy to the local firewall, which enforces it.

Enter the following commands to set the appropriate client firewall parameters. You can configure only one instance of this command. [Table 30-2](#), following this set of commands, explains the syntax elements of these commands:

Cisco Integrated Firewall

```
hostname(config-group-policy)# client-firewall {opt | req} cisco-integrated acl-in ACL
acl-out ACL
```

Cisco Security Agent

```
hostname(config-group-policy)# client-firewall {opt | req} cisco-security-agent
```

No Firewall

```
hostname(config-group-policy)# client-firewall none
```

Custom Firewall

```
hostname(config-group-policy)# client-firewall {opt | req} custom vendor-id num product-id
num policy {AYT | CPP acl-in ACL acl-out ACL} [description string]
```

Zone Labs Firewalls

```
hostname(config-group-policy)# client-firewall {opt | req} zonelabs-integrity
```



Note

When the firewall type is **zonelabs-integrity**, do not include arguments. The Zone Labs Integrity Server determines the policies.

```
hostname(config-group-policy)# client-firewall {opt | req} zonelabs-zonealarm policy {AYT
| CPP acl-in ACL acl-out ACL}
```

```
hostname(config-group-policy)# client-firewall {opt | req} zonelabs-zonealarmpro policy
{AYT | CPP acl-in ACL acl-out ACL}
```

```
client-firewall {opt | req} zonelabs-zonealarmpro policy {AYT | CPP acl-in ACL acl-out
ACL}
```

Sygate Personal Firewalls

```
hostname(config-group-policy)# client-firewall {opt | req} sygate-personal

hostname(config-group-policy)# client-firewall {opt | req} sygate-personal-pro

hostname(config-group-policy)# client-firewall {opt | req} sygate-security-agent
```

Network Ice, Black Ice Firewall:

```
hostname(config-group-policy)# client-firewall {opt | req} networkkice-blackkice
```

Table 30-2 *client-firewall Command Keywords and Variables*

Parameter	Description
acl-in <i>ACL</i>	Provides the policy the client uses for inbound traffic.
acl-out <i>ACL</i>	Provides the policy the client uses for outbound traffic.
AYT	Specifies that the client PC firewall application controls the firewall policy. The security appliance checks to make sure that the firewall is running. It asks, “Are You There?” If there is no response, the security appliance tears down the tunnel.
cisco-integrated	Specifies Cisco Integrated firewall type.
cisco-security-agent	Specifies Cisco Intrusion Prevention Security Agent firewall type.
CPP	Specifies Policy Pushed as source of the VPN client firewall policy.
custom	Specifies Custom firewall type.
description <i>string</i>	Describes the firewall.
networkkice-blackkice	Specifies Network ICE Black ICE firewall type.
none	Indicates that there is no client firewall policy. Sets a firewall policy with a null value, thereby disallowing a firewall policy. Prevents inheriting a firewall policy from a default or specified group policy.
opt	Indicates an optional firewall type.
product-id	Identifies the firewall product.
req	Indicates a required firewall type.
sygate-personal	Specifies Sygate Personal firewall type.
sygate-personal-pro	Specifies Sygate Personal Pro firewall type.
sygate-security-agent	Specifies Sygate Security Agent firewall type.
vendor-id	Identifies the firewall vendor.
zonelabs-integrity	Specifies Zone Labs Integrity Server firewall type.
zonelabs-zonealarm	Specifies Zone Labs Zone Alarm firewall type.
zonelabs-zonealarmorpro policy	Specifies Zone Labs Zone Alarm or Pro firewall type.
zonelabs-zonealarmpro policy	Specifies Zone Labs Zone Alarm Pro firewall type.

The following example shows how to set a client firewall policy that requires Cisco Intrusion Prevention Security Agent for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# client-firewall req cisco-security-agent
hostname(config-group-policy)#
```

Configuring Client Access Rules

Configure rules that limit the remote access client types and versions that can connect via IPSec through the security appliance by using the **client-access-rule** command in group-policy configuration mode. Construct rules according to these guidelines:

- If you do not define any rules, the security appliance permits all connection types.
- When a client matches none of the rules, the security appliance denies the connection. If you define a deny rule, you must also define at least one permit rule; otherwise, the security appliance denies all connections.
- For both software and hardware clients, type and version must exactly match their appearance in the **show vpn-sessiondb remote** display.
- The * character is a wildcard, which you can enter multiple times in each rule. For example, **client-access rule 3 deny type * version 3.*** creates a priority 3 client access rule that denies all client types running release versions 3.x software.
- You can construct a maximum of 25 rules per group policy.
- There is a limit of 255 characters for an entire set of rules.
- You can enter n/a for clients that do not send client type and/or version.

To delete a rule, enter the **no** form of this command. This command is equivalent to the following command:

```
hostname(config-group-policy)# client-access-rule 1 deny type "Cisco VPN Client" version 4.0
```

To delete all rules, enter the **no client-access-rule** command without arguments. This deletes all configured rules, including a null rule if you created one by issuing the **client-access-rule** command with the **none** keyword.

By default, there are no access rules. When there are no client access rules, users inherit any rules that exist in the default group policy.

To prevent users from inheriting client access rules, enter the **client-access-rule** command with the **none** keyword. The result of this command is that all client types and versions can connect.

```
hostname(config-group-policy)# client-access rule priority {permit | deny} type type
version {version | none}
```

```
hostname(config-group-policy)# no client-access rule [priority {permit | deny} type type
version version]
```

Table 30-3 explains the meaning of the keywords and parameters in these commands.

Table 30-3 *client-access rule Command Keywords and Variables*

Parameter	Description
deny	Denies connections for devices of a particular type and/or version.
none	Allows no client access rules. Sets client-access-rule to a null value, thereby allowing no restriction. Prevents inheriting a value from a default or specified group policy.

Table 30-3 *client-access rule Command Keywords and Variables*

permit	Permits connections for devices of a particular type and/or version.
<i>priority</i>	Determines the priority of the rule. The rule with the lowest integer has the highest priority. Therefore, the rule with the lowest integer that matches a client type and/or version is the rule that applies. If a lower priority rule contradicts, the security appliance ignores it.
type <i>type</i>	Identifies device types via free-form strings, for example VPN 3002. A string must match exactly its appearance in the show vpn-sessiondb remote display, except that you can enter the * character as a wildcard.
version <i>version</i>	Identifies the device version via free-form strings, for example 7.0. A string must match exactly its appearance in the show vpn-sessiondb remote display, except that you can enter the * character as a wildcard.

The following example shows how to create client access rules for the group policy named FirstGroup. These rules permit Cisco VPN clients running software version 4.x, while denying all Windows NT clients:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# client-access-rule 1 deny type WinNT version *
hostname(config-group-policy)# client-access-rule 2 permit "Cisco VPN Client" version 4.*
```



Note The “type” field is a free-form string that allows any value, but that value must match the fixed value that the client sends to the security appliance at connect time.

Configuring Group-Policy WebVPN Attributes

WebVPN lets users establish a secure, remote-access VPN tunnel to the security appliance using a web browser. There is no need for either a software or hardware client. WebVPN provides easy access to a broad range of web resources and web-enabled applications from almost any computer that can reach HTTPS Internet sites. WebVPN uses SSL and its successor, TLS1, to provide a secure connection between remote users and specific, supported internal resources that you configure at a central site. The security appliance recognizes connections that need to be proxied, and the HTTP server interacts with the authentication subsystem to authenticate users. By default, WebVPN is disabled.

You can customize a WebVPN configuration for specific internal group policies.



Note

The webvpn mode that you enter from global configuration mode lets you configure global settings for WebVPN. The webvpn mode described in this section, which you enter from group-policy configuration mode, lets you customize a WebVPN configuration for specific group policies.

In group-policy webvpn configuration mode, you can specify whether to inherit the settings for all the functions or customize the following parameters, each of which is described in the subsequent sections:

- functions
- customizations
- html-content-filter
- homepage
- filter

- url-list
- port-forward
- port-forward-name
- sso server (single-signon server)
- auto-signon
- deny message
- SSL VPN Client (SVC)
- keep-alive ignore
- HTTP compression

In many instances, you define the WebVPN attributes as part of configuring WebVPN, then you apply those definitions to specific groups when you configure the group-policy webvpn attributes. Enter group-policy webvpn configuration mode by using the **webvpn** command in group-policy configuration mode. WebVPN commands for group policies define access to files, MAPI proxy, URLs and TCP applications over WebVPN. They also identify ACLs and types of traffic to filter. WebVPN is disabled by default. See the description of WebVPN in *Cisco Security Appliance Command Line Configuration Guide* and *Cisco Security Appliance Command Reference* for more information about configuring the WebVPN attributes.

To remove all commands entered in group-policy webvpn configuration mode, enter the **no** form of this command. These webvpn commands apply to the username or group policy from which you configure them.

```
hostname(config-group-policy)# webvpn
hostname(config-group-policy)# no webvpn
```

You do not need to configure WebVPN to use e-mail proxies.

The following example shows how to enter group-policy webvpn configuration mode for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)#
```

Configuring Group-Policy WebVPN Function Attributes

Configure the WebVPN functions that you want to enable. To configure file access and file browsing, HTTP Proxy, MAPI Proxy, and URL entry over WebVPN for this group policy, enter the **functions** command in webvpn mode. These functions are disabled by default.

```
hostname(config-group-webvpn)# functions {auto-download | citrix | file-access |  
file-browsing | file-entry | filter | http-proxy | mapi | none | port-forward | url-entry}
```

To remove a configured function, enter the **no** form of this command. The **no** option allows inheritance of a value from another group policy. To prevent inheriting function values, enter the **functions none** command.

To remove all configured functions, including a null value created by issuing the **functions none** command, enter the **no** form of this command without arguments.

```
hostname(config-group-webvpn)# no functions [auto-download | citrix | file-access |  
file-browsing | file-entry | filter | http-proxy | mapi | none | port-forward | url-entry]
```

Table 30-4 describes the meaning of the keywords used in this command.

Table 30-4 *functions Command Keywords*

Keyword	Meaning
auto-download	Automatically downloads the port forwarding applet
citrix	Enables Citrix
file-access	Enables or disables file access. When enabled, the WebVPN home page lists file servers in the server list. You must enable file access to enable file browsing and/or file entry.
file-browsing	Enables or disables browsing for file servers and shares. You must enable file browsing to allow user entry of a file server.
file-entry	Enables or disables user ability to enter names of file servers.
filter	Applies a webtype access-list
http-proxy	Enables or disables the forwarding of an HTTP applet proxy to the client. The proxy is useful for technologies that interfere with proper mangling, such as Java, ActiveX, and Flash. It bypasses mangling while ensuring the continued use of the security appliance. The forwarded proxy modifies the browser's old proxy configuration automatically and redirects all HTTP and HTTPS requests to the new proxy configuration. It supports virtually all client side technologies, including HTML, CSS, JavaScript, VBScript, ActiveX, and Java. The only browser it supports is Microsoft Internet Explorer.
mapi	Enables or disables Microsoft Outlook/Exchange port forwarding.
none	Sets a null value for all WebVPN functions . Prevents inheriting functions from a default or specified group policy
port-forward	Enables port forwarding
url-entry	Enables or disables user entry of URLs. When enabled, the security appliance still restricts URLs with any configured URL or network ACLs. When URL entry is disabled, the security appliance restricts WebVPN users to the URLs on the home page.

The following example shows how to configure file access, file browsing, and MAPI Proxy for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# functions file-access file-browsing MAPI
hostname(config-group-webvpn)#
```

Applying Customization

Customizations determine the appearance of the windows that the user sees upon login. You configure the customization parameters as part of configuring WebVPN. To apply a previously defined web-page customization to change the look-and-feel of the web page that the user sees at login, enter the customization command in group-policy webvpn configuration mode:

```
hostname(config-group-webvpn)# customization customization_name
hostname(config-group-webvpn)#
```

For example, to use the customization named blueborder, enter the following command:

```
hostname(config-group-webvpn)# customization blueborder
```

```
hostname(config-group-webvpn)#
```

You configure the customization itself by entering the **customization** command in WebVPN mode.

The following example shows a command sequence that first establishes a WebVPN customization named 123 that defines a password prompt. The example then defines a WebVPN group policy named testpolicy and uses the **customization** command to specify the use of the WebVPN customization named 123:

```
hostname(config)# webvpn
hostname(config-webvpn)# customization 123
hostname(config-webvpn-custom)# password-prompt Enter password
hostname(config-webvpn)# exit
hostname(config)# group-policy testpolicy nopassword
hostname(config)# group-policy testpolicy attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# customization value 123
hostname(config-group-webvpn)#
```

Specifying a “Deny” Message

You can specify the message delivered to a remote user who logs into WebVPN successfully, but has no VPN privileges by entering the **deny-message** command in group-policy webvpn configuration mode:

```
hostname(config-group-webvpn)# deny-message value "message"
hostname(config-group-webvpn)# no deny-message value "message"
hostname(config-group-webvpn)# deny-message none
```

The **no deny-message value** command removes the message string, so that the remote user does not receive a message.

The **no deny-message none** command removes the attribute from the tunnel group policy configuration. The policy inherits the attribute value.

The message can be up to 491 alphanumeric characters long, including special characters, spaces, and punctuation, but not counting the enclosing quotation marks. The text appears on the remote user’s browser upon login. When typing the string in the **deny-message value** command, continue typing even if the command wraps.

The default deny message is: “Login was successful, but because certain criteria have not been met or due to some specific group policy, you do not have permission to use any of the VPN features. Contact your IT administrator for more information.”

The first command in the following example creates an internal group policy named group2. The subsequent commands modify the attributes, including the webvpn deny message associated with that policy.

```
hostname(config)# group-policy group2 internal
hostname(config)# group-policy group2 attributes
hostname(config-group)# webvpn
hostname(config-group-webvpn)# deny-message value "Your login credentials are OK. However,
you have not been granted rights to use the VPN features. Contact your administrator for
more information."
hostname(config-group-webvpn)
```

Configuring Group-Policy WebVPN Filter Attributes

Specify whether to filter Java, ActiveX, images, scripts, and cookies for WebVPN sessions for this group policy by using the **html-content-filter** command in webvpn mode. HTML filtering is disabled by default.

To remove a content filter, enter the **no** form of this command. To remove all content filters, including a null value created by issuing the **html-content-filter** command with the **none** keyword, enter the **no** form of this command without arguments. The **no** option allows inheritance of a value from another group policy. To prevent inheriting an html content filter, enter the **html-content-filter** command with the **none** keyword.

Using the command a second time overrides the previous setting.

```
hostname(config-group-webvpn)# html-content-filter {java | images | scripts | cookies | none}
```

```
hostname(config-group-webvpn)# no html-content-filter [java | images | scripts | cookies | none]
```

Table 30-5 describes the meaning of the keywords used in this command.

Table 30-5 *filter Command Keywords*

Keyword	Meaning
cookies	Removes cookies from images, providing limited ad filtering and privacy.
images	Removes references to images (removes tags).
java	Removes references to Java and ActiveX (removes <EMBED>, <APPLET>, and <OBJECT> tags).
none	Indicates that there is no filtering. Sets a null value, thereby disallowing filtering. Prevents inheriting filtering values.
scripts	Removes references to scripting (removes <SCRIPT> tags).

The following example shows how to set filtering of JAVA and ActiveX, cookies, and images for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# html-content-filter java cookies images
hostname(config-group-webvpn)#
```

Specifying the User Home Page

Specify a URL for the web page that displays when a user in this group logs in by using the **homepage** command in group-policy webvpn configuration mode. There is no default home page.

To remove a configured home page, including a null value created by issuing the **homepage none** command, enter the **no** form of this command. The **no** option allows inheritance of a value from another group policy. To prevent inheriting a home page, enter the **homepage none** command.

The **none** keyword indicates that there is no WebVPN home page. It sets a null value, thereby disallowing a home page and prevents inheriting an home page.

The *url-string* variable following the keyword **value** provides a URL for the home page. The string must begin with either http:// or https://.

```
hostname(config-group-webvpn)# homepage {value url-string | none}
hostname(config-group-webvpn)# no homepage
hostname(config-group-webvpn)#
```

Configuring Auto-Signon

The **auto-signon** command is a single sign-on method for WebVPN users. It passes the WebVPN login credentials (username and password) to internal servers for authentication using NTLM authentication, basic authentication, or both. Multiple auto-signon commands can be entered and are processed according to the input order (early commands take precedence).

You can use the auto-signon feature in three modes: webvpn configuration, webvpn group configuration, or webvpn username configuration mode. The typical precedence behavior applies where username supersedes group, and group supersedes global. The mode you choose depends upon the desired scope of authentication.

To disable auto-signon for a particular user to a particular server, use the **no** form of the command with the original specification of IP block or URI. To disable authentication to all servers, use the **no** form without arguments. The **no** option allows inheritance of a value from the group policy.

The following example, entered in group-policy webvpn configuration mode, configures auto-signon for the user named anyuser, using basic authentication, to servers with IP addresses ranging from 10.1.1.0 to 10.1.1.255:

The following example commands configure auto-signon for WebVPN users ExamplePolicy group policy, using either basic or NTLM authentication, to servers defined by the URI mask `https://*.example.com/*`:

```
hostname(config)# group-policy ExamplePolicy attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# auto-signon allow uri https://*.example.com/* auth-type all
hostname(config-group-webvpn)#
```

The following example commands configure auto-signon for WebVPN users ExamplePolicy2 group policy, using either basic or NTLM authentication, to the server with the IP address 10.1.1.0, using subnet mask 255.255.255.0:

```
hostname(config)# group-policy ExamplePolicy attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# auto-signon allow ip 10.1.1.0 255.255.255.0 auth-type all
hostname(config-group-webvpn)#
```

Specifying the Access List to Use for WebVPN Connections

Specify the name of the access list to use for WebVPN connections for this group policy or username by using the **filter** command in webvpn mode. WebVPN access lists do not apply until you enter the **filter** command to specify them.

To remove the access list, including a null value created by issuing the **filter none** command, enter the **no** form of this command. The **no** option allows inheritance of a value from another group policy. To prevent inheriting filter values, enter the **filter value none** command.

WebVPN access lists do not apply until you enter the **filter** command to specify them.

You configure ACLs to permit or deny various types of traffic for this group policy. You then enter the **filter** command to apply those ACLs for WebVPN traffic.

```
hostname(config-group-webvpn)# filter {value ACLname | none}
hostname(config-group-webvpn)# no filter
```

The **none** keyword indicates that there is no **webvpntype** access list. It sets a null value, thereby disallowing an access list and prevents inheriting an access list from another group policy.

The *ACLname* string following the keyword **value** provides the name of the previously configured access list.

**Note**

You must have **filter** selected in the **functions** command. WebVPN does not use ACLs defined in the **vpn-filter** command.

The following example shows how to set a filter that invokes an access list named *acl_in* for the group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# filter acl_in
hostname(config-group-webvpn)#
```

Applying a URL List

You can specify a list of URLs to appear on the WebVPN home page for a group policy. First, you must create one or more named lists by entering the **url-list** command in global configuration mode. To apply a list of WebVPN servers and URLs to a particular group policy, allowing access to the URLs in a list for a specific group policy, use the name of the list or lists you create there with the **url-list** command in group-policy webvpn configuration mode. There is no default URL list.

To remove a list, including a null value created by using the **url-list none** command, use the **no** form of this command. The **no** option allows inheritance of a value from another group policy. To prevent inheriting a URL list, use the **url-list none** command. Using the command a second time overrides the previous setting:

```
hostname(config-group-webvpn)# url-list {value name | none} [index]
hostname(config-group-webvpn)# no url-list
```

Table 30-6 shows the **url-list** command parameters and their meanings.

Table 30-6 *url-list Command Keywords and Variables*

Parameter	Meaning
<i>index</i>	Indicates the display priority on the home page.
none	Sets a null value for url lists. Prevents inheriting a list from a default or specified group policy.
<i>value name</i>	Specifies the name of a previously configured list of urls. To configure such a list, use the url-list command in global configuration mode.

The following example sets a URL list called FirstGroupURLs for the group policy named FirstGroup and specifies that this should be the first URL list displayed on the homepage:

```
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# url-list value FirstGroupURLs 1
hostname(config-group-webvpn)#
```

Enabling WebVPN Application Access for a Group Policy

To enable WebVPN application access for this group policy, enter the **port-forward** command in group-policy webvpn configuration mode. Port forwarding is disabled by default.

Before you can enter the **port-forward** command in group-policy webvpn configuration mode to enable application access, you must define a list of applications that you want users to be able to use in a WebVPN connection. Enter the **port-forward** command in global configuration mode to define this list.

To remove the port forwarding attribute from the group-policy configuration, including a null value created by issuing the **port-forward none** command, enter the **no** form of this command. The **no** option allows inheritance of a list from another group policy. To prevent inheriting a port forwarding list, enter the **port-forward** command with the **none** keyword. The **none** keyword indicates that there is no filtering. It sets a null value, thereby disallowing a filtering, and prevents inheriting filtering values.

The syntax of the command is as follows:

```
hostname(config-group-webvpn)# port-forward {value listname | none}
hostname(config-group-webvpn)# no port-forward
```

The *listname* string following the keyword **value** identifies the list of applications WebVPN users can access. Enter the port-forward command in webvpn configuration mode to define the list.

Using the command a second time overrides the previous setting.

The following example shows how to set a port-forwarding list called *ports1* for the internal group policy named *FirstGroup*:

```
hostname(config)# group-policy FirstGroup internal attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# port-forward value ports1
hostname(config-group-webvpn)#
```

Configuring the Port-Forwarding Display Name

Configure the display name that identifies TCP port forwarding to end users for a particular user or group policy by using the **port-forward-name** command in group-policy webvpn configuration mode. To delete the display name, including a null value created by using the **port-forward-name none** command, enter the **no** form of the command. The **no** option restores the default name, Application Access. To prevent a display name, enter the **port-forward none** command. The syntax of the command is as follows:

```
hostname(config-group-webvpn)# port-forward-name {value name | none}
hostname(config-group-webvpn)# no port-forward-name
```

The following example shows how to set the name, Remote Access TCP Applications, for the internal group policy named *FirstGroup*:

```
hostname(config)# group-policy FirstGroup internal attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# port-forward-name value Remote Access TCP Applications
hostname(config-group-webvpn)#
```

Configuring the Maximum Object Size to Ignore for Updating the Session Timer

Network devices exchange short keepalive messages to ensure that the virtual circuit between them is still active. The length of these messages can vary. The **keep-alive-ignore** command lets you tell the security appliance to consider all messages that are less than or equal to the specified size as keepalive messages and not as traffic when updating the session timer. The range is 0 through 900 KB. The default is 4 KB.

To specify the upper limit of the HTTP/HTTPS traffic, per transaction, to ignore, use the **keep-alive-ignore** command in group-policy attributes webvpn configuration mode:

```
hostname(config-group-webvpn)# keep-alive-ignore size
hostname(config-group-webvpn)#
```

The **no** form of the command removes this specification from the configuration:

```
hostname(config-group-webvpn)# no keep-alive-ignore
hostname(config-group-webvpn)#
```

The following example sets the maximum size of objects to ignore as 5 KB:

```
hostname(config-group-webvpn) # keep-alive-ignore 5
hostname(config-group-webvpn) #
```

Specifying HTTP Compression

Enable compression of http data over a WebVPN connection for a specific group or user by entering the **http-comp** command in the group policy webvpn mode.

```
hostname(config-group-webvpn) # http-comp {gzip | none}
hostname(config-group-webvpn) #
```

To remove the command from the configuration and cause the value to be inherited, use the **no** form of the command:

```
hostname(config-group-webvpn) # no http-comp {gzip | none}
hostname(config-group-webvpn) #
```

The syntax of this command is as follows:

- **gzip**—Specifies compression is enabled for the group or user. This is the default value.
- **none**—Specifies compression is disabled for the group or user.

For WebVPN connections, the **compression** command configured from global configuration mode overrides the **http-comp** command configured in group policy and username webvpn modes.

In the following example, compression is disabled for the group-policy sales:

```
hostname(config) # group-policy sales attributes
hostname(config-group-policy) # webvpn
hostname(config-group-webvpn) # http-comp none
hostname(config-group-webvpn) #
```

Specifying the SSO Server

Single sign-on support, available only for WebVPN, lets users access different secure services on different servers without reentering a username and password more than once. The **sso-server value** command, when entered in group-policy-webvpn mode, lets you assign an SSO server to a group policy.

To assign an SSO server to a group policy, use the **sso-server value** command in group-policy-webvpn configuration mode. This command requires that your configuration include CA SiteMinder command.

```
hostname(config-group-webvpn) # sso-server value server_name
hostname(config-group-webvpn) #
```

To remove the assignment and use the default policy, use the **no** form of this command. To prevent inheriting the default policy, use the **sso-server none** command.

```
hostname(config-group-webvpn) # sso-server {value server_name | none}
hostname(config-group-webvpn) # [no] sso-server value server_name
```

The default policy assigned to the SSO server is DfltGrpPolicy.

The following example creates the group policy “my-sso-grp-pol” and assigns it to the SSO server named “example”:

```
hostname(config) # group-policy my-sso-grp-pol internal
hostname(config) # group-policy my-sso-grp-pol attributes
hostname(config-group-policy) # webvpn
hostname(config-group-webvpn) # sso-server value example
hostname(config-group-webvpn) #
```

Configuring SVC

The SSL VPN Client (SVC) is a VPN tunneling technology that gives remote users the benefits of an IPsec VPN client without the need for network administrators to install and configure IPsec VPN clients on remote computers. The SVC uses the SSL encryption that is already present on the remote computer as well as the WebVPN login and authentication of the security appliance.

To establish an SVC session, the remote user enters the IP address of a WebVPN interface of the security appliance in the browser, and the browser connects to that interface and displays the WebVPN login screen. If the user satisfies the login and authentication, and the security appliance identifies the user as *requiring* the SVC, the security appliance downloads the SVC to the remote computer. If the security appliance identifies the user as having the *option* to use the SVC, the security appliance downloads the SVC to the remote computer while presenting a link on the user screen to skip the SVC installation.

After downloading, the SVC installs and configures itself, and then the SVC either remains or uninstalls itself (depending on the configuration) from the remote computer when the connection terminates.

The security appliance might have several unique SVC images residing in cache memory for different remote computer operating systems. When the user attempts to connect, the security appliance can consecutively download portions of these images to the remote computer until the image and operating system match, at which point it downloads the entire SVC. You can order the SVC images to minimize connection setup time, with the first image downloaded representing the most commonly-encountered remote computer operating system. For complete information about installing and using SVC, see [Chapter 38, “Configuring SSL VPN Client”](#).

After enabling SVC, as described in [Chapter 38, “Configuring SSL VPN Client”](#), you can enable or require SVC features for a specific group. This feature is disabled by default. If you enable or require SVC, you can then enable a succession of svc commands, described in this section. To enable SVC and its related svc commands, do the following steps in group-policy webvpn configuration mode:

- Step 1** To enable the security appliance to download SVC files to remote computers, enter the **svc enable** command. By default, this command is disabled. The security appliance does not download SVC files. To remove the **svc enable** command from the configuration, use the **no** form of this command.

```
hostname(config-group-webvpn)# svc {none | enable | required}
hostname(config-group-webvpn)#
```



Note

Entering the **no svc enable** command does not terminate active SVC sessions.

```
hostname(config)# group-policy sales attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# svc enable
hostname(config-group-webvpn)#
```

- Step 2** To enable compression of HTTP data over an SVC connection, for a specific group, enter the svc compression command. By default, SVC compression is set to **deflate** (enabled). To disable compression for a specific group, use the **none** keyword. To remove the svc compression command and cause the value to be inherited, use the **no** form of the command:

```
hostname(config-group-webvpn)# svc compression {deflate | none}
hostname(config-group-webvpn)#
```

The following example disables SVC compression for the group policy named sales:

```
hostname(config)# group-policy sales attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# svc compression none
hostname(config-group-webvpn)#
```


- Step 3** To enable dead-peer-detection (DPD) on the security appliance and to set the frequency with which either the SVC or the security appliance performs DPD, use the **svc dpd-interval** command. To remove the **svc dpd-interval** command from the configuration, use the **no** form of the command. To disable SVC DPD for this group, use the **none** keyword:

```
hostname(config-group-webvpn)# svc dpd-interval {[gateway {seconds | none}] | [client {seconds | none}]}
```

DPD checking is disabled by default.

The gateway refers to the security appliance. You can specify the frequency with which the security appliance performs the DPD test as a range of from 30 to 3600 seconds (1 hour). Specifying **none** disables the DPD testing that the security appliance performs.

The client refers to the SVC. You can specify the frequency with which the client performs the DPD test as a range of from 30 to 3600 seconds (1 hour). Specifying **none** disables the DPD testing that the client performs.

In the following example, the user configures the DPD frequency performed by the security appliance (gateway) to 3000 seconds, and the DPD frequency performed by the client to 1000 seconds for the existing group policy named sales:

```
hostname(config)# group-policy sales attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# svc dpd-interval gateway 3000
hostname(config-group-webvpn)# svc dpd-interval client 1000
hostname(config-group-webvpn)#
```

- Step 4** You can adjust the frequency of keepalive messages (specified by *seconds*), to ensure that an SVC connection through a proxy, firewall, or NAT device remains open, even if the device limits the time that the connection can be idle.

Adjusting the frequency also ensures that the SVC does not disconnect and reconnect when the remote user is not actively running a socket-based application, such as Microsoft Outlook or Microsoft Internet Explorer.

To configure the frequency (15 through 600 seconds) which an SVC on a remote computer sends keepalive messages to the security appliance, use the **svc keepalive** command. Use the **no** form of the command to remove the command from the configuration and cause the value to be inherited:

```
hostname(config-group-webvpn)# svc keepalive {none | seconds}
hostname(config-group-webvpn)# no svc keepalive {none | seconds}
```

SVC keepalives are disabled by default. Using the keyword **none** disables SVC keepalive messages.

The following example configures the security appliance to enable the SVC to send keepalive messages, with a frequency of 300 seconds (5 minutes):

```
hostname(config-group-webvpn)# svc keepalive 300
hostname(config-group-webvpn)#
```

- Step 5** To enable the permanent installation of an SVC onto a remote computer, use the **svc keep-installer** command with the **installed** keyword. To remove the command from the configuration, use the **no** form of this command:

```
hostname(config-group-webvpn)# svc keep-installer {installed | none}
hostname(config-group-webvpn)# no svc keep-installer {installed | none}
```

The default is that permanent installation of the SVC is disabled. The SVC uninstalls from the remote computer at the end of the SVC session.

The following example configures the security appliance to keep the SVC installed on the remote computer for this group:

```
hostname(config-group-webvpn)# svc keep-installer installed
hostname(config-group-webvpn)#
```

Step 6 To enable the SVC to perform a rekey on an SVC session, use the **svc rekey** command. To disable rekey and remove the command from the configuration, use the **no** form of this command:

```
hostname(config-group-webvpn)# svc rekey {method {ssl | new-tunnel} | time minutes | none}}
hostname(config-group-webvpn)# no svc rekey {method {ssl | new-tunnel} | time minutes | none}}
hostname(config-group-webvpn)#
```

By default, SVC rekey is disabled.

Specifying the method as new-tunnel specifies that the SVC establishes a new tunnel during SVC rekey. Specifying the method as none disables SVC rekey. Specifying the method as ssl specifies that SSL renegotiation takes place during SVC rekey. Instead of specifying the method, you can specify the time; that is, the number of minutes from the start of the session until the re-key takes place, from 1 through 10080 (1 week).

For the **no** form of the command, only the minimum is necessary, as the following example shows:

```
hostname(config-username-webvpn)# no svc rekey method
hostname(config-username-webvpn)#
```

If, however, you specify the method as new-tunnel:

```
hostname(config-username-webvpn)# no svc rekey method new-tunnel
hostname(config-username-webvpn)#
```

but the current method is ssl, then the command fails, because the values don't match.

In the following example, the user configures the SVC to renegotiate with SSL during rekey and configures the rekey to occur 30 minutes after the session begins:

```
hostname(config-group-webvpn)# svc rekey method ssl
hostname(config-group-webvpn)# svc rekey time 30
hostname(config-group-webvpn)#
```

Configuring User Attributes

This section describes user attributes and how to configure them. It includes the following sections:

- [Viewing the Username Configuration, page 30-71](#)
- [Configuring Attributes for Specific Users, page 30-71](#)

By default, users inherit all user attributes from the assigned group policy. The security appliance also lets you assign individual attributes at the user level, overriding values in the group policy that applies to that user. For example, you can specify a group policy giving all users access during business hours, but give a specific user 24-hour access.

Viewing the Username Configuration

To display the configuration for all usernames, including default values inherited from the group policy, enter the **all** keyword with the **show running-config username** command, as follows:

```
hostname# show running-config all username
hostname#
```

This displays the encrypted password and the privilege level. for all users, or, if you supply a username, for that specific user. If you omit the **all** keyword, only explicitly configured values appear in this list. The following example displays the output of this command for the user named testuser:

```
hostname# show running-config all username testuser
username testuser password 12RsxXQnphyr/I9Z encrypted privilege 15
```

Configuring Attributes for Specific Users

To configure specific users, you assign a password (or no password) and attributes to a user using the **username** command, which enters username mode. Any attributes that you do not specify are inherited from the group policy.

The internal user authentication database consists of the users entered with the **username** command. The **login** command uses this database for authentication. To add a user to the security appliance database, enter the **username** command in global configuration mode. To remove a user, use the **no** version of this command with the username you want to remove. To remove all usernames, use the **clear configure username** command without appending a username.

Setting a User Password and Privilege Level

Enter the **username** command to assign a password and a privilege level for a user. You can enter the **nopassword** keyword to specify that this user does not require a password. If you do specify a password, you can specify whether that password is stored in an encrypted form.

The optional **privilege** keyword lets you set a privilege level for this user. Privilege levels range from 0 (the lowest) through 15. System administrators generally have the highest privilege level. The default level is 2.

```
hostname(config)# username name {nopassword | password password [encrypted]} [privilege
priv_level]}
```

```
hostname(config)# no username [name]
```

Table 30-7 describes the meaning of the keywords and variables used in this command.

Table 30-7 *username Command Keywords and Variables*

Keyword/Variable	Meaning
encrypted	Indicates that the password is encrypted.
<i>name</i>	Provides the name of the user.
nopassword	Indicates that this user needs no password.

password <i>password</i>	Indicates that this user has a password, and provides the password.
privilege <i>priv_level</i>	Sets a privilege level for this user. The range is from 0 to 15, with lower numbers having less ability to use commands and administer the security appliance. The default privilege level is 2. The typical privilege level for a system administrator is 15.

By default, VPN users that you add with this command have no attributes or group policy association. You must explicitly configure all values.

The following example shows how to configure a user named anyuser with an encrypted password of pw_12345678 and a privilege level of 12:

```
hostname(config)# username anyuser password pw_12345678 encrypted privilege 12
hostname(config)#
```

Configuring User Attributes

After configuring the user's password (if any) and privilege level, you set the other attributes. These can be in any order. To remove any attribute-value pair, enter the **no** form of the command.

Enter username mode by entering the **username** command with the **attributes** keyword:

```
hostname(config)# username name attributes
hostname(config-username)#
```

The prompt changes to indicate the new mode. You can now configure the attributes.

Configuring VPN User Attributes

The VPN user attributes set values specific to VPN connections, as described in the following sections.

Configuring Inheritance

You can let users inherit from the group policy the values of attributes that you have not configured at the username level. To specify the name of the group policy from which this user inherits attributes, enter the **vpn-group-policy** command. By default, VPN users have no group-policy association:

```
hostname(config-username)# vpn-group-policy group-policy-name
hostname(config-username)# no vpn-group-policy group-policy-name
```

For an attribute that is available in username mode, you can override the value of an attribute in a group policy for a particular user by configuring it in username mode.

The following example shows how to configure a user named anyuser to use attributes from the group policy named FirstGroup:

```
hostname(config)# username anyuser attributes
hostname(config-username)# vpn-group-policy FirstGroup
hostname(config-username)#
```

Configuring Access Hours

Associate the hours that this user is allowed to access the system by specifying the name of a configured time-range policy:

To remove the attribute from the running configuration, enter the **no** form of this command. This option allows inheritance of a time-range value from another group policy. To prevent inheriting a value, enter the **vpn-access-hours none** command. The default is unrestricted access.

```
hostname(config-username)# vpn-access-hours value {time-range | none}
hostname(config-username)# vpn-access-hours value none
hostname(config)#
```

The following example shows how to associate the user named anyuser with a time-range policy called 824:

```
hostname(config)# username anyuser attributes
hostname(config-username)# vpn-access-hours 824
hostname(config-username)#
```

Configuring Maximum Simultaneous Logins

Specify the maximum number of simultaneous logins allowed for this user. The range is 0 through 2147483647. The default is 3 simultaneous logins. To remove the attribute from the running configuration, enter the **no** form of this command. Enter 0 to disable login and prevent user access.

```
hostname(config-username)# vpn-simultaneous-logins integer
hostname(config-username)# no vpn-simultaneous-logins
hostname(config-username)#
```



Note

While the maximum limit for the number of simultaneous logins is very large, allowing several could compromise security and affect performance.

The following example shows how to allow a maximum of 4 simultaneous logins for the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# vpn-simultaneous-logins 4
hostname(config-username)#
```

Configuring the Idle Timeout

Specify the idle timeout period in minutes, or enter **none** to disable the idle timeout. If there is no communication activity on the connection in this period, the security appliance terminates the connection.

The range is 1 through 35791394 minutes. The default is 30 minutes. To allow an unlimited timeout period, and thus prevent inheriting a timeout value, enter the **vpn-idle-timeout** command with the **none** keyword. To remove the attribute from the running configuration, enter the **no** form of this command.

```
hostname(config-username)# vpn-idle-timeout {minutes | none}
hostname(config-username)# no vpn-idle-timeout
hostname(config-username)#
```

The following example shows how to set a VPN idle timeout of 15 minutes for the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# vpn-idle-timeout 30
hostname(config-username)#
```

Configuring the Maximum Connect Time

Specify the maximum user connection time in minutes, or enter **none** to allow unlimited connection time and prevent inheriting a value for this attribute. At the end of this period of time, the security appliance terminates the connection.

The range is 1 through 35791394 minutes. There is no default timeout. To allow an unlimited timeout period, and thus prevent inheriting a timeout value, enter the **vpn-session-timeout** command with the **none** keyword. To remove the attribute from the running configuration, enter the **no** form of this command.

```
hostname(config-username)# vpn-session-timeout {minutes | none}
hostname(config-username)# no vpn-session-timeout
hostname(config-username)#
```

The following example shows how to set a VPN session timeout of 180 minutes for the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# vpn-session-timeout 180
hostname(config-username)#
```

Applying an ACL Filter

Specify the name of a previously-configured, user-specific ACL to use as a filter for VPN connections. To disallow an access list and prevent inheriting an access list from the group policy, enter the **vpn-filter** command with the **none** keyword. To remove the ACL, including a null value created by issuing the **vpn-filter none** command, enter the **no** form of this command. The **no** option allows inheritance of a value from the group policy. There are no default behaviors or values for this command.

You configure ACLs to permit or deny various types of traffic for this user. You then use the **vpn-filter** command to apply those ACLs.

```
hostname(config-username)# vpn-filter {value ACL_name | none}
hostname(config-username)# no vpn-filter
hostname(config-username)#
```



Note

You must have **filter** selected in the **functions** command. WebVPN does not use ACLs defined in the **vpn-filter** command.

The following example shows how to set a filter that invokes an access list named `acl_vpn` for the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# vpn-filter value acl_vpn
hostname(config-username)#
```

Specifying the IP Address and Netmask

Specify the IP address and netmask to assign to a particular user. To remove the IP address, enter the **no** form of this command.

```
hostname(config-username)# vpn-framed-ip-address {ip_address}
hostname(config-username)# no vpn-framed-ip-address
hostname(config-username)#
```

The following example shows how to set an IP address of 10.92.166.7 for a user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# vpn-framed-ip-address 10.92.166.7
hostname(config-username)
```

Specify the network mask to use with the IP address specified in the previous step. If you used the **no vpn-framed-ip-address** command, do not specify a network mask. To remove the subnet mask, enter the **no** form of this command. There is no default behavior or value.

```
hostname(config-username)# vpn-framed-ip-netmask {netmask}
hostname(config-username)# no vpn-framed-ip-netmask
hostname(config-username)
```

The following example shows how to set a subnet mask of 255.255.255. 254 for a user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# vpn-framed-ip-netmask 255.255.255.254
hostname(config-username)
```

Specifying the Tunnel Protocol

Specify the VPN tunnel types (IPSec or WebVPN) that this user can use. The default is taken from the default group policy, the default for which is IPSec. To remove the attribute from the running configuration, enter the **no** form of this command.

```
hostname(config-username)# vpn-tunnel-protocol {webvpn | IPSec}
hostname(config-username)# no vpn-tunnel-protocol [webvpn | IPSec]
hostname(config-username)
```

The parameter values for this command are as follows:

- **IPSec**—Negotiates an IPSec tunnel between two peers (a remote access client or another secure gateway). Creates security associations that govern authentication, encryption, encapsulation, and key management.
- **webvpn**—Provides VPN services to remote users via an HTTPS-enabled web browser, and does not require a client

Enter this command to configure one or more tunneling modes. You must configure at least one tunneling mode for users to connect over a VPN tunnel.

The following example shows how to configure WebVPN and IPSec tunneling modes for the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# vpn-tunnel-protocol webvpn
hostname(config-username)# vpn-tunnel-protocol IPSec
hostname(config-username)
```

Restricting Remote User Access

Configure the **group-lock** attribute with the **value** keyword to restrict remote users to access only through the specified, preexisting tunnel group. Group-lock restricts users by checking whether the group configured in the VPN client is the same as the tunnel group to which the user is assigned. If it is not, the security appliance prevents the user from connecting. If you do not configure group-lock, the security appliance authenticates users without regard to the assigned group.

To remove the **group-lock** attribute from the running configuration, enter the **no** form of this command. This option allows inheritance of a value from the group policy. To disable group-lock, and to prevent inheriting a group-lock value from a default or specified group policy, enter the **group-lock** command with the **none** keyword.

```
hostname(config-username)# group-lock {value tunnel-grp-name | none}
hostname(config-username)# no group-lock
hostname(config-username)
```

The following example shows how to set group lock for the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# group-lock value tunnel-group-name
hostname(config-username)
```

Enabling Password Storage for Software Client Users

Specify whether to let users store their login passwords on the client system. Password storage is disabled by default. Enable password storage only on systems that you know to be in secure sites. To disable password storage, enter the **password-storage** command with the **disable** keyword. To remove the password-storage attribute from the running configuration, enter the **no** form of this command. This enables inheritance of a value for password-storage from the group policy.

```
hostname(config-username)# password-storage {enable | disable}
hostname(config-username)# no password-storage
hostname(config-username)
```

This command has no bearing on interactive hardware client authentication or individual user authentication for hardware clients.

The following example shows how to enable password storage for the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# password-storage enable
hostname(config-username)
```

Configuring WebVPN for Specific Users

The following sections describe how to customize a WebVPN configuration for specific users. Enter username webvpn configuration mode by using the **webvpn** command in username configuration mode. WebVPN lets users establish a secure, remote-access VPN tunnel to the security appliance using a web browser. There is no need for either a software or hardware client. WebVPN provides easy access to a broad range of web resources and web-enabled applications from almost any computer that can reach HTTPS Internet sites. WebVPN uses SSL and its successor, TLS1, to provide a secure connection between remote users and specific, supported internal resources that you configure at a central site. The security appliance recognizes connections that need to be proxied, and the HTTP server interacts with the authentication subsystem to authenticate users.

The username webvpn configuration mode commands define access to files, MAPI proxy, URLs and TCP applications over WebVPN. They also identify ACLs and types of traffic to filter. WebVPN is disabled by default. These **webvpn** commands apply only to the username from which you configure them. Notice that the prompt changes, indicating that you are now in username webvpn configuration mode.

```
hostname(config-username)# webvpn
hostname(config-username-webvpn)#
```

To remove all commands entered in username webvpn configuration mode, use the **no** form of this command:

```
hostname(config-username)# no webvpn
hostname(config-username)#
```


You do not need to configure WebVPN to use e-mail proxies.



Note

The webvpn mode that you enter from global configuration mode lets you configure global settings for WebVPN. The username webvpn configuration mode described in this section, which you enter from username mode, lets you customize a WebVPN configuration for specific users.

In username webvpn configuration mode, you can customize the following parameters, each of which is described in the subsequent steps:

- functions
- customizations
- deny message
- html-content-filter
- homepage
- filter
- url-list
- port-forward
- port-forward-name
- sso server (single-signon server)
- auto-signon
- SSL VPN Client (SVC)
- keep-alive ignore
- HTTP compression

The following example shows how to enter username webvpn configuration mode for the username anyuser attributes:

```
hostname(config)# username anyuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)#
```

Configuring the WebVPN Functions to Enable

To configure the WebVPN functions file access and file browsing, HTTP Proxy, MAPI Proxy, and URL entry over WebVPN for this user, enter the **functions** command in username webvpn configuration mode. To remove a configured function, enter the **no** form of this command. These functions are disabled by default.

To remove all configured functions, including a null value created by issuing the **functions none** command, enter the **no** form of this command without arguments. The **no** option allows inheritance of a value from another group policy. To prevent inheriting function values, enter the **functions none** command.

```
hostname(config-username-webvpn)# functions {auto-download | citrix | file-access |
file-browsing | file-entry | filter | http-proxy | mapi | none | port-forward | url-entry}

hostname(config-username-webvpn)# no functions [auto-download | citrix | file-access |
file-browsing | file-entry | filter | http-proxy | mapi | none | port-forward | url-entry]
```

Table 30-8 describes the meaning of the keywords used in this command.

Table 30-8 *functions Command Keywords*

Keyword	Meaning
auto-download	Automatically downloads the port forwarding applet
citrix	Enables Citrix
file-access	Enables or disables file access. When enabled, the WebVPN home page lists file servers in the server list. You must enable file access to enable file browsing and/or file entry.
file-browsing	Enables or disables browsing for file servers and shares. You must enable file browsing to allow user entry of a file server.
file-entry	Enables or disables user ability to enter names of file servers.
filter	Applies a webtype access-list
http-proxy	Enables or disables the forwarding of an HTTP applet proxy to the client. The proxy is useful for technologies that interfere with proper mangling, such as Java, ActiveX, and Flash. It bypasses mangling while ensuring the continued use of the security appliance. The forwarded proxy modifies the browser's old proxy configuration automatically and redirects all HTTP and HTTPS requests to the new proxy configuration. It supports virtually all client side technologies, including HTML, CSS, JavaScript, VBScript, ActiveX, and Java. The only browser it supports is Microsoft Internet Explorer.
mapi	Enables or disables Microsoft Outlook/Exchange port forwarding.
none	Sets a null value for all WebVPN functions . Prevents inheriting functions from a default or specified group policy
port-forward	Enables port forwarding
url-entry	Enables or disables user entry of URLs. When enabled, the security appliance still restricts URLs with any configured URL or network ACLs. When URL entry is disabled, the security appliance restricts WebVPN users to the URLs on the home page.

The keywords used in this command are as follows:

- **file-access**—Enables or disables file access. When enabled, the WebVPN home page lists file servers in the server list. You must enable file access to enable file browsing and/or file entry.
- **file-browsing**—Enables or disables browsing for file servers and shares. You must enable file browsing to allow user entry of a file server.
- **file-entry**—Enables or disables user ability to enter names of file servers.
- **http-proxy**—Enables or disables the forwarding of an HTTP applet proxy to the client. The proxy is useful for technologies that interfere with proper mangling, such as Java, ActiveX, and Flash. It bypasses mangling while ensuring the continued use of the security appliance. The forwarded proxy modifies the browser's old proxy configuration automatically and redirects all HTTP and HTTPS requests to the new proxy configuration. It supports virtually all client side technologies, including HTML, CSS, JavaScript, VBScript, ActiveX, and Java. The only browser it supports is Microsoft Internet Explorer.
- **mapi**—Enables or disables Microsoft Outlook/Exchange port forwarding.

- **none**—Sets a null value for all WebVPN **functions**. Prevents inheriting functions from a default or specified group policy
- **url-entry**—Enables or disables user entry of URLs. When enabled, the security appliance still restricts URLs with any configured URL or network ACLs. When URL entry is disabled, the security appliance restricts WebVPN users to the URLs on the home page.

The following example shows how to configure file access, file browsing, HTTP Proxy, and MAPI Proxy for the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# functions file-access file-browsing MAPI
hostname(config-username-webvpn)#
```

Specifying the Content/Objects to Filter from the HTML

To filter Java, ActiveX, images, scripts, and cookies for WebVPN sessions for this user, enter the **html-content-filter** command in username webvpn configuration mode. To remove a content filter, enter the **no** form of this command. To remove all content filters, including a null value created by issuing the **html-content-filter none** command, enter the **no** form of this command without arguments. The **no** option allows inheritance of a value from the group policy. To prevent inheriting an HTML content filter, enter the **html-content-filter none** command. HTML filtering is disabled by default.

Using the command a second time overrides the previous setting.

```
hostname(config-username-webvpn)# html-content-filter {java | images | scripts | cookies | none}

hostname(config-username-webvpn)# no html-content-filter [java | images | scripts | cookies | none]
```

The keywords used in this command are as follows:

- **cookies**—Removes cookies from images, providing limited ad filtering and privacy.
- **images**—Removes references to images (removes tags).
- **java**—Removes references to Java and ActiveX (removes <EMBED>, <APPLET>, and <OBJECT> tags).
- **none**—Indicates that there is no filtering. Sets a null value, thereby disallowing filtering. Prevents inheriting filtering values.
- **scripts**—Removes references to scripting (removes <SCRIPT> tags).

The following example shows how to set filtering of JAVA and ActiveX, cookies, and images for the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# html-content-filter java cookies images
hostname(config-username-webvpn)#
```

Specifying the User Home Page

To specify a URL for the web page that displays when this WebVPN user logs in, enter the **homepage** command in username webvpn configuration mode. To remove a configured home page, including a null value created by issuing the **homepage none** command, enter the **no** form of this command. The **no** option allows inheritance of a value from the group policy. To prevent inheriting a home page, enter the **homepage none** command.

The **none** keyword indicates that there is no WebVPN home page. It sets a null value, thereby disallowing a home page and prevents inheriting a home page.

The *url-string* variable following the keyword **value** provides a URL for the home page. The string must begin with either `http://` or `https://`.

There is no default home page.

```
hostname(config-username-webvpn)# homepage {value url-string | none}
hostname(config-username-webvpn)# no homepage
hostname(config-username-webvpn)#
```

The following example shows how to specify `www.example.com` as the home page for the user named `anyuser`:

```
hostname(config)# username anyuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# homepage value www.example.com
hostname(config-username-webvpn)#
```

Applying Customization

Customizations determine the appearance of the windows that the user sees upon login. You configure the customization parameters as part of configuring WebVPN. To apply a previously defined web-page customization to change the look-and-feel of the web page that the user sees at login, enter the customization command in username webvpn configuration mode:

```
hostname(config-username-webvpn)# customization {none | value customization_name}
hostname(config-username-webvpn)#
```

For example, to use the customization named `blueborder`, enter the following command:

```
hostname(config-username-webvpn)# customization value blueborder
hostname(config-username-webvpn)#
```

You configure the customization itself by entering the **customization** command in WebVPN mode.

The following example shows a command sequence that first establishes a WebVPN customization named `123` that defines a password prompt. The example then defines a WebVPN tunnel-group named `test` and uses the **customization** command to specifies the use of the WebVPN customization named `123`:

```
hostname(config)# webvpn
hostname(config-webvpn)# customization 123
hostname(config-webvpn-custom)# password-prompt Enter password
hostname(config-webvpn)# exit
hostname(config)# username testuser nopassword
hostname(config)# username testuser attributes
hostname(config-username-webvpn)# webvpn
hostname(config-username-webvpn)# customization value 123
hostname(config-username-webvpn)#
```

Specifying a “Deny” Message

You can specify the message delivered to a remote user who logs into WebVPN successfully, but has no VPN privileges by entering the **deny-message** command in username webvpn configuration mode:

```
hostname(config-username-webvpn)# deny-message value "message"
hostname(config-username-webvpn)# no deny-message value "message"
hostname(config-username-webvpn)# deny-message none
```

The **no deny-message value** command removes the message string, so that the remote user does not receive a message.

The **no deny-message none** command removes the attribute from the tunnel group policy configuration. The policy inherits the attribute value.

The message can be up to 491 alphanumeric characters long, including special characters, spaces, and punctuation, but not counting the enclosing quotation marks. The text appears on the remote user's browser upon login. When typing the string in the **deny-message value** command, continue typing even if the command wraps.

The default deny message is: “Login was successful, but because certain criteria have not been met or due to some specific group policy, you do not have permission to use any of the VPN features. Contact your IT administrator for more information.”

The first command in the following example enters username mode and configures the attributes for the user named anyuser. The subsequent commands enter username webvpn configuration mode and modify the deny message associated with that user.

```
hostname(config)# username anyuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# deny-message value "Your login credentials are OK. However, you have not been granted rights to use the VPN features. Contact your administrator for more information."
hostname(config-username-webvpn)
```

Specifying the Access List to use for WebVPN Connections

To specify the name of the access list to use for WebVPN connections for this user, enter the **filter** command in username webvpn configuration mode. To remove the access list, including a null value created by issuing the **filter none** command, enter the **no** form of this command. The **no** option allows inheritance of a value from the group policy. To prevent inheriting filter values, enter the **filter value none** command.

WebVPN access lists do not apply until you enter the **filter** command to specify them.

You configure ACLs to permit or deny various types of traffic for this user. You then enter the **filter** command to apply those ACLs for WebVPN traffic.

```
hostname(config-username-webvpn)# filter {value ACLname | none}
hostname(config-username-webvpn)# no filter
hostname(config-username-webvpn)#
```

The **none** keyword indicates that there is no **webvpntype** access list. It sets a null value, thereby disallowing an access list and prevents inheriting an access list from another group policy.

The *ACLname* string following the keyword **value** provides the name of the previously configured access list.



Note

WebVPN does not use ACLs defined in the **vpn-filter** command.

The following example shows how to set a filter that invokes an access list named *acl_in* for the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# filter acl_in
hostname(config-username-webvpn)#
```

Applying a URL List

You can specify a list of URLs to appear on the WebVPN home page for a user. First, you must create one or more named lists by entering the **url-list** command in global configuration mode. To apply a list of WebVPN servers and URLs to a particular user, enter the **url-list** command in username webvpn configuration mode.

To remove a list, including a null value created by using the **url-list none** command, enter the **no** form of this command. The **no** option allows inheritance of a value from the group policy. To prevent inheriting a url list, enter the **url-list none** command.

```
hostname(config-username-webvpn)# url-list {listname displayname url | none}
hostname(config-username-webvpn)# no url-list
```

The keywords and variables used in this command are as follows:

- *displayname*—Specifies a name for the URL. This name appears on the WebVPN end user interface.
- *listname*—Identifies a name by which to group URLs.
- **none**—Indicates that there is no list of URLs. Sets a null value, thereby disallowing a URL list. Prevents inheriting URL list values.
- *url*—Specifies a URL that WebVPN users can access.

There is no default URL list.

Using the command a second time overrides the previous setting.

The following example shows how to set a URL list called AnyuserURLs for the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# url-list value AnyuserURLs
hostname(config-username-webvpn)#
```

Enabling WebVPN Application Access

To enable WebVPN application access for this user, enter the **port-forward** command in username webvpn configuration mode. Port forwarding is disabled by default.

To remove the port forwarding attribute from the configuration, including a null value created by issuing the **port-forward none** command, enter the **no** form of this command. The **no** option allows inheritance of a list from the group policy. To disallow filtering and prevent inheriting a port forwarding list, enter the **port-forward** command with the **none** keyword.

```
hostname(config-username-webvpn)# port-forward {value listname | none}
hostname(config-username-webvpn)# no port-forward
hostname(config-username-webvpn)#
```

The *listname* string following the keyword **value** identifies the list of applications WebVPN users can access. Enter the **port-forward** command in configuration mode to define the list.

Using the command a second time overrides the previous setting.

Before you can enter the **port-forward** command in username webvpn configuration mode to enable application access, you must define a list of applications that you want users to be able to use in a WebVPN connection. Enter the **port-forward** command in global configuration mode to define this list.

The following example shows how to configure a portforwarding list called ports1:

```
hostname(config-group-policy)# webvpn
hostname(config-username-webvpn)# port-forward value ports1
hostname(config-username-webvpn)#
```

Configuring the Port-Forwarding Display Name

Configure the display name that identifies TCP port forwarding to end users for a particular user by using the **port-forward-name** command in username webvpn configuration mode. To delete the display name, including a null value created by using the **port-forward-name none** command, enter the **no** form of the command. The **no** option restores the default name, Application Access. To prevent a display name, enter the **port-forward none** command.

```
hostname(config-username-webvpn)# port-forward-name {value name | none}
hostname(config-username-webvpn)# no port-forward-name
```

The following example shows how to configure the port-forward name test:

```
hostname(config-group-policy)# webvpn
hostname(config-username-webvpn)# port-forward-name value test
hostname(config-username-webvpn)#
```

Configuring the Maximum Object Size to Ignore for Updating the Session Timer

Network devices exchange short keepalive messages to ensure that the virtual circuit between them is still active. The length of these messages can vary. The **keep-alive-ignore** command lets you tell the security appliance to consider all messages that are less than or equal to the specified size as keepalive messages and not as traffic when updating the session timer. The range is 0 through 900 KB. The default is 4 KB.

To specify the upper limit of the HTTP/HTTPS traffic, per transaction, to ignore, use the **keep-alive-ignore** command in group-policy attributes webvpn configuration mode:

```
hostname(config-group-webvpn)# keep-alive-ignore size
hostname(config-group-webvpn)#
```

The **no** form of the command removes this specification from the configuration:

```
hostname(config-group-webvpn)# no keep-alive-ignore
hostname(config-group-webvpn)#
```

The following example sets the maximum size of objects to ignore as 5 KB:

```
hostname(config-group-webvpn)# keep-alive-ignore 5
hostname(config-group-webvpn)#
```

Configuring Auto-Signon

To automatically submit the WebVPN login credentials of a particular WebVPN user to internal servers using NTLM, basic HTTP authentication or both, use the **auto-signon** command in username webvpn configuration mode.

The **auto-signon** command is a single sign-on method for WebVPN users. It passes the WebVPN login credentials (username and password) to internal servers for authentication using NTLM authentication, basic authentication, or both. Multiple auto-signon commands can be entered and are processed according to the input order (early commands take precedence).

You can use the auto-signon feature in three modes: webvpn configuration, webvpn group configuration, or webvpn username configuration mode. The typical precedence behavior applies where username supersedes group, and group supersedes global. The mode you choose will depend upon the desired scope of authentication.

To disable auto-signon for a particular user to a particular server, use the **no** form of the command with the original specification of IP block or URI. To disable authentication to all servers, use the **no** form without arguments. The **no** option allows inheritance of a value from the group policy.

The following example commands configure auto-signon for a WebVPN user named anyuser, using either basic or NTLM authentication, to servers defined by the URI mask `https://*.example.com/*`:

```
hostname(config)# username anyuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# auto-signon allow uri https://*.example.com/* auth-type
all
```

The following example commands configure auto-signon for the WebVPN user named anyuser, using either basic or NTLM authentication, to the server with the IP address 10.1.1.0, using subnet mask 255.255.255.0:

```
hostname(config)# username anyuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# auto-signon allow ip 10.1.1.0 255.255.255.0 auth-type
all
hostname(config-username-webvpn)#
```

Specifying HTTP Compression

Enable compression of http data over a WebVPN connection for a specific user by entering the **http-comp** command in the username webvpn configuration mode.

```
hostname(config-username-webvpn)# http-comp {gzip | none}
hostname(config-username-webvpn)#
```

To remove the command from the configuration and cause the value to be inherited, use the **no** form of the command:

```
hostname(config-username-webvpn)# no http-comp {gzip | none}
hostname(config-username-webvpn)#
```

The syntax of this command is as follows:

- **gzip**—Specifies compression is enabled for the group or user. This is the default value.
- **none**—Specifies compression is disabled for the group or user.

For WebVPN connections, the **compression** command configured from global configuration mode overrides the **http-comp** command configured in group policy and username webvpn modes.

In the following example, compression is disabled for the username testuser:

```
hostname(config)# username testuser internal
hostname(config)# username testuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# http-comp none
hostname(config-username-webvpn)#
```


Specifying the SSO Server

Single sign-on support, available only for WebVPN, lets users access different secure services on different servers without reentering a username and password more than once. The **sso-server value** command, when entered in username-webvpn mode, lets you assign an SSO server to a user.

To assign an SSO server to a user, use the **sso-server value** command in username-webvpn configuration mode. This command requires that your configuration include CA SiteMinder command.

```
hostname(config-username-webvpn)# sso-server value server_name
hostname(config-username-webvpn)#
```

To remove the assignment and use the default policy, use the **no** form of this command. To prevent inheriting the default policy, use the **sso-server none** command.

```
hostname(config-username-webvpn)# sso-server {value server_name | none}
hostname(config-username-webvpn)# [no] sso-server value server_name
```

The default policy assigned to the SSO server is DfltGrpPolicy.

The following example assigns the SSO server named example to the user named anyuser:

```
hostname(config)# username anyuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# sso-server value example
hostname(config-username-webvpn)#
```

Configuring SVC

The SSL VPN Client (SVC) is a VPN tunneling technology that gives remote users the benefits of an IPSec VPN client without the need for network administrators to install and configure IPSec VPN clients on remote computers. The SVC uses the SSL encryption that is already present on the remote computer as well as the WebVPN login and authentication of the security appliance.

To establish an SVC session, the remote user enters the IP address of a WebVPN interface of the security appliance in the browser, and the browser connects to that interface and displays the WebVPN login screen. If the user satisfies the login and authentication, and the security appliance identifies the user as *requiring* the SVC, the security appliance downloads the SVC to the remote computer. If the security appliance identifies the user as having the *option* to use the SVC, the security appliance downloads the SVC to the remote computer while presenting a link on the user screen to skip the SVC installation.

After downloading, the SVC installs and configures itself, and then the SVC either remains or uninstalls itself (depending on the configuration) from the remote computer when the connection terminates.

The security appliance might have several unique SVC images residing in cache memory for different remote computer operating systems. When the user attempts to connect, the security appliance can consecutively download portions of these images to the remote computer until the image and operating system match, at which point it downloads the entire SVC. You can order the SVC images to minimize connection setup time, with the first image downloaded representing the most commonly-encountered remote computer operating system. For complete information about installing and using SVC, see [Chapter 38, “Configuring SSL VPN Client”](#).

After enabling SVC, as described in [Chapter 38, “Configuring SSL VPN Client”](#), you can enable or require SVC features for a specific user. This feature is disabled by default. If you enable or require SVC, you can then enable a succession of svc commands, described in this section. To enable SVC and its related svc commands, do the following steps in username webvpn configuration mode:

- Step 1** To enable the security appliance to download SVC files to remote computers, enter the **svc enable** command. By default, this command is disabled. The security appliance does not download SVC files. To remove the **svc enable** command from the configuration, use the **no** form of this command.

```
hostname(config-username-webvpn)# svc {none | enable | required}
hostname(config-username-webvpn)#
```

**Note**

Entering the **no svc enable** command does not terminate active SVC sessions.

```
hostname(config)# username sales attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# svc enable
hostname(config-username-webvpn)#
```

- Step 2** To enable compression of HTTP data over an SVC connection, for a specific user, enter the **svc compression** command. By default, SVC compression is set to **deflate** (enabled). To disable compression for a specific user, use the **none** keyword. To remove the **svc compression** command and cause the value to be inherited, use the **no** form of the command:

```
hostname(config-username-webvpn)# svc compression {deflate | none}
hostname(config-username-webvpn)#
```

The following example disables SVC compression for the user named sales:

```
hostname(config)# username sales attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# svc compression none
hostname(config-username-webvpn)#
```

- Step 3** To enable dead-peer-detection (DPD) on the security appliance and to set the frequency with which either the SVC or the security appliance performs DPD, use the **svc dpd-interval** command. To remove the **svc dpd-interval** command from the configuration, use the **no** form of the command. To disable SVC DPD for this user, use the **none** keyword:

```
hostname(config-username-webvpn)# svc dpd-interval {[gateway {seconds | none}] | [client {seconds | none}]}
hostname(config-username-webvpn)#
```

DPD checking is disabled by default.

The gateway refers to the security appliance. You can specify the frequency with which the security appliance performs the DPD test as a range of from 30 to 3600 seconds (1 hour). Specifying **none** disables the DPD testing that the security appliance performs.

The client refers to the SVC. You can specify the frequency with which the client performs the DPD test as a range of from 30 to 3600 seconds (1 hour). Specifying **none** disables the DPD testing that the client performs.

In the following example, the user configures the DPD frequency performed by the security appliance (gateway) to 3000 seconds, and the DPD frequency performed by the client to 1000 seconds for the existing user named sales:

```
hostname(config)# username sales attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# svc dpd-interval gateway 3000
hostname(config-username-webvpn)# svc dpd-interval client 1000
hostname(config-username-webvpn)#
```

- Step 4** You can adjust the frequency of keepalive messages (specified by *seconds*), to ensure that an SVC connection through a proxy, firewall, or NAT device remains open, even if the device limits the time that the connection can be idle.

Adjusting the frequency also ensures that the SVC does not disconnect and reconnect when the remote user is not actively running a socket-based application, such as Microsoft Outlook or Microsoft Internet Explorer.

To configure the frequency (15 through 600 seconds) which an SVC on a remote computer sends keepalive messages to the security appliance, use the **svc keepalive** command. Use the **no** form of the command to remove the command from the configuration and cause the value to be inherited:

```
hostname(config-username-webvpn) # svc keepalive {none | seconds}
hostname(config-username-webvpn) # no svc keepalive {none | seconds}
hostname(config-username-webvpn) #
```

SVC keepalives are disabled by default. Using the keyword **none** disables SVC keepalive messages.

In the following example, the user configures the security appliance to enable the SVC to send keepalive messages, with a frequency of 300 seconds (5 minutes):

```
hostname(config-username-webvpn) # svc keepalive 300
hostname(config-username-webvpn) #
```

- Step 5** To enable the permanent installation of an SVC onto a remote computer, use the **svc keep-installer** command with the **installed** keyword. To remove the command from the configuration, use the **no** form of this command:

```
hostname(config-username-webvpn) # svc keep-installer {installed | none}
hostname(config-username-webvpn) # no svc keep-installer {installed | none}
hostname(config-username-webvpn) #
```

The default is that permanent installation of the SVC is disabled. The SVC uninstalls from the remote computer at the end of the SVC session.

The following example configures the security appliance to keep the SVC installed on the remote computer for this user:

```
hostname(config-username-webvpn) # svc keep-installer installed
hostname(config-username-webvpn) #
```

- Step 6** To enable the SVC to perform a rekey on an SVC session, use the **svc rekey** command:

```
hostname(config-username-webvpn) # svc rekey {method {ssl | new-tunnel} | time minutes | none}}
```

To disable rekey and remove the command from the configuration, use the **no** form of this command:

```
hostname(config-username-webvpn) # no svc rekey [method {ssl | new-tunnel} | time minutes | none}}
hostname(config-username-webvpn) #
```

By default, SVC rekey is disabled.

Specifying the method as **new-tunnel** specifies that the SVC establishes a new tunnel during SVC rekey. Specifying the method as **none** disables SVC rekey. Specifying the method as **ssl** specifies that SSL renegotiation takes place during SVC rekey. Instead of specifying the method, you can specify the time; that is, the number of minutes from the start of the session until the re-key takes place, from 1 through 10080 (1 week).

For the **no** form of the command, only the minimum is necessary. The following example is correct:

```
hostname(config-username-webvpn) # no svc rekey method
hostname(config-username-webvpn) #
```

If, however, you specify the method as new-tunnel:

```
hostname(config-username-webvpn)# no svc rekey method new-tunnel
hostname(config-username-webvpn)#
```

and the current method is ssl, then the command fails, because the values don't match.

In the following example, the user configures the SVC to renegotiate with SSL during rekey and configures the rekey to occur 30 minutes after the session begins:

```
hostname(config-username-webvpn)# svc rekey method ssl
hostname(config-username-webvpn)# svc rekey time 30
hostname(config-username-webvpn)#
```



CHAPTER 31

Configuring IP Addresses for VPNs

This chapter describes IP address assignment methods.

IP addresses make internetwork connections possible. They are like telephone numbers: both the sender and receiver must have an assigned number to connect. But with VPNs, there are actually two sets of addresses: the first set connects client and server on the public network. Once that connection is made, the second set connects client and server through the VPN tunnel.

In security appliance address management, we are dealing with the second set of IP addresses: those private IP addresses that connect a client with a resource on the private network, through the tunnel, and let the client function as if it were directly connected to the private network. Furthermore, we are dealing only with the private IP addresses that get assigned to clients. The IP addresses assigned to other resources on your private network are part of your network administration responsibilities, not part of VPN management. Therefore, when we discuss IP addresses here, we mean those IP addresses available in your private network addressing scheme that let the client function as a tunnel endpoint.

This chapter includes the following sections:

- [Configuring an IP Address Assignment Method, page 31-1](#)
- [Configuring Local IP Address Pools, page 31-2](#)
- [Configuring AAA Addressing, page 31-2](#)
- [Configuring DHCP Addressing, page 31-3](#)

Configuring an IP Address Assignment Method

The security appliance can use one or more of the following methods for assigning IP addresses to remote access clients. If you configure more than one address assignment method, the security appliance searches each of the options until it finds an IP address. By default, all methods are enabled. To view the current configuration, enter the **show running-config all vpn-addr-assign** command.

- **aaa**—Retrieves addresses from an external authentication server on a per-user basis. If you are using an authentication server that has IP addresses configured, we recommend using this method.
- **dhcp**—Obtains IP addresses from a DHCP server. If you want to use DHCP, you must configure a DHCP server. You must also define the range of IP addresses that the DHCP server can use.
- **local**—Use an internal address pool. Internally configured address pools are the easiest method of address pool assignment to configure. If you choose local, you must also use the **ip-local-pool** command to define the range of IP addresses to use.

To specify a method for assigning IP addresses to remote access clients, enter the **vpn-addr-assign** command in global configuration mode. The syntax is **vpn-addr-assign {aaa | dhcp | local}**.

Configuring Local IP Address Pools

To configure IP address pools to use for VPN remote access tunnels, enter the **ip local pool** command in global configuration mode. To delete address pools, enter the **no** form of this command.

The security appliance uses address pools based on the tunnel group for the connection. If you configure more than one address pool for a tunnel group, the security appliance uses them in the order in which they are configured.

If you assign addresses from a non-local subnet, we suggest that you add pools that fall on subnet boundaries to make adding routes for these networks easier.

A summary of the configuration of local address pools follows:

```
hostname(config)# vpn-addr-assign local
hostname(config)# ip local pool firstpool 10.20.30.40-10.20.30.50 mask 255.255.255.0
hostname(config)
```

-
- Step 1** To configure IP address pools as the address assignment method, enter the **vpn-addr-assign** command with the **local** argument:

```
hostname(config)# vpn-addr-assign local
hostname(config)#
```

- Step 2** To configure an address pool, enter the **ip local pool** command. The syntax is **ip local pool poolname first-address—last-address mask mask**.

The following example configures an IP address pool named firstpool. The starting address is 10.20.30.40 and the ending address is 10.20.30.50. The network mask is 255.255.255.0.

```
hostname(config)# ip local pool firstpool 10.20.30.40-10.20.30.50 mask 255.255.255.0
hostname(config)
```

Configuring AAA Addressing

To use a AAA server to assign addresses for VPN remote access clients, you must first configure a AAA server or server group. See the **aaa-server protocol** command in the *Cisco Security Appliance Command Reference* and “[Identifying AAA Server Groups and Servers](#),” in [Chapter 13, “Configuring AAA Servers and the Local Database”](#) of this guide.

In addition, the user must match a tunnel group configured for RADIUS authentication.

The following examples illustrate how to define a AAA server group called RAD2 for the tunnel group named firstgroup. It includes one more step than is necessary, in that previously you might have named the tunnel group and defined the tunnel group type. This step appears in the following example as a reminder that you have no access to subsequent tunnel-group commands until you set these values.

An overview of the configuration that these examples create follows:

```
hostname(config)# vpn-addr-assign aaa
hostname(config)# tunnel-group firstgroup type ipsec-ra
hostname(config)# tunnel-group firstgroup general-attributes
hostname(config-general)# authentication-server-group RAD2
```

To configure AAA for IP addressing, perform the following steps:

- Step 1** To configure AAA as the address assignment method, enter the **vpn-addr-assign** command with the **aaa** argument:

```
hostname(config)# vpn-addr-assign aaa
hostname(config)#
```

- Step 2** To establish the tunnel group called firstgroup as a remote access or LAN-to-LAN tunnel group, enter the **tunnel-group** command with the **type** keyword. The following example configures a remote access tunnel group.

```
hostname(config)# tunnel-group firstgroup type ipsec-ra
hostname(config)#
```

- Step 3** To enter general-attributes configuration mode, which lets you define a AAA server group for the tunnel group called firstgroup, enter the **tunnel-group** command with the **general-attributes** argument.

```
hostname(config)# tunnel-group firstgroup general-attributes
hostname(config-general)#
```

- Step 4** To specify the AAA server group to use for authentication, enter the **authentication-server-group** command.

```
hostname(config-general)# authentication-server-group RAD2
hostname(config-general)#
```

This command has more arguments that this example includes. For more information, see the *Cisco Security Appliance Command Reference*.

Configuring DHCP Addressing

To use DHCP to assign addresses for VPN clients, you must first configure a DHCP server and the range of IP addresses that the DHCP server can use. Then you define the DHCP server on a tunnel group basis. Optionally, you can also define a DHCP network scope in the group policy associated with the tunnel group or username. This is either an IP network number or IP Address that identifies to the DHCP server which pool of IP addresses to use.

The following examples define the DHCP server at IP address 172.33.44.19 for the tunnel group named firstgroup. They also define a DHCP network scope of 192.86.0.0 for the group policy called remotegroup. (The group policy called remotegroup is associated with the tunnel group called firstgroup). If you do not define a network scope, the DHCP server assigns IP addresses in the order of the address pools configured. It goes through the pools until it identifies an unassigned address.

The following configuration includes more steps than are necessary, in that previously you might have named and defined the tunnel group type as remote access, and named and identified the group policy as internal or external. These steps appear in the following examples as a reminder that you have no access to subsequent tunnel-group and group-policy commands until you set these values.

A summary of the configuration that these examples create follows:

```
hostname(config)# vpn-addr-assign dhcp
hostname(config)# tunnel-group firstgroup type ipsec-ra
hostname(config)# tunnel-group firstgroup general-attributes
hostname(config-general)# dhcp-server 172.33.44.19
hostname(config-general)# exit
hostname(config)# group-policy remotegroup internal
hostname(config)# group-policy remotegroup attributes
hostname(config-group-policy)# dhcp-network-scope 192.86.0.0
```

To define a DHCP server for IP addressing, perform the following steps.

-
- Step 1** To configure DHCP as the address assignment method, enter the **vpn-addr-assign** command with the **dhcp** argument:
- ```
hostname(config)# vpn-addr-assign dhcp
hostname(config)#
```
- Step 2** To establish the tunnel group called firstgroup as a remote access or LAN-to-LAN tunnel group, enter the **tunnel-group** command with the **type** keyword. The following example configures a remote access tunnel group.
- ```
hostname(config)# tunnel-group firstgroup type ipsec-ra
hostname(config)#
```
- Step 3** To enter general-attributes configuration mode, which lets you configure a DHCP server, enter the **tunnel-group** command with the **general-attributes** argument.
- ```
hostname(config)# tunnel-group firstgroup general-attributes
hostname(config)#
```
- Step 4** To define the DHCP server, enter the **dhcp-server** command. The following example configures a DHCP server at IP address 172.33.44.19.
- ```
hostname(config-general)# dhcp-server 172.33.44.19
hostname(config-general)#
```
- Step 5** Exit tunnel-group mode.
- ```
hostname(config-general)# exit
hostname(config)#
```
- Step 6** To define the group policy called remotegroup as an internally or externally configured group, enter the **group-policy** command with the **internal** or **external** argument. The following example configures an internal group.
- ```
hostname(config)# group-policy remotegroup internal
hostname(config)#
```
- Step 7** (Optional) To enter group-policy attributes configuration mode, which lets you configure a subnetwork of IP addresses for the DHCP server to use, enter the **group-policy** command with the **attributes** keyword.
- ```
hostname(config)# group-policy remotegroup attributes
hostname(config-group-policy)#
```
- Step 8** (Optional) To specify the range of IP addresses the DHCP server should use to assign addresses to users of the group policy called remotegroup, enter the **dhcp-network-scope** command. The following example configures at network scope of 192.86.0.0.
- ```
hostname(config-group-policy)# dhcp-network-scope 192.86.0.0
hostname(config-group-policy)#
```
-



CHAPTER 32

Configuring Remote Access IPsec VPNs

Remote access VPNs let single users connect to a central site through a secure connection over a TCP/IP network such as the Internet.

This chapter describes how to build a remote access VPN connection. It includes the following sections:

- [Summary of the Configuration, page 32-1](#)
- [Configuring Interfaces, page 32-2](#)
- [Configuring ISAKMP Policy and Enabling ISAKMP on the Outside Interface, page 32-3](#)
- [Configuring an Address Pool, page 32-4](#)
- [Adding a User, page 32-4](#)
- [Creating a Transform Set, page 32-4](#)
- [Defining a Tunnel Group, page 32-5](#)
- [Creating a Dynamic Crypto Map, page 32-6](#)
- [Creating a Crypto Map Entry to Use the Dynamic Crypto Map, page 32-7](#)

Summary of the Configuration

This chapter uses the following configuration to explain how to configure a remote access connection. Later sections provide step-by-step instructions.

```
hostname(config)# interface ethernet0
hostname(config-if)# ip address 10.10.4.200 255.255.0.0
hostname(config-if)# nameif outside
hostname(config)# no shutdown
hostname(config)# isakmp policy 1 authentication pre-share
hostname(config)# isakmp policy 1 encryption 3des
hostname(config)# isakmp policy 1 hash sha
hostname(config)# isakmp policy 1 group 2
hostname(config)# isakmp policy 1 lifetime 43200
hostname(config)# isakmp enable outside
hostname(config)# ip local pool testpool 192.168.0.10-192.168.0.15
hostname(config)# username testuser password 12345678
hostname(config)# crypto ipsec transform set FirstSet esp-3des esp-md5-hmac
hostname(config)# tunnel-group testgroup type ipsec-ra
hostname(config)# tunnel-group testgroup general-attributes
hostname(config-general)# address-pool testpool
hostname(config)# tunnel-group testgroup ipsec-attributes
hostname(config-ipsec)# pre-shared-key 44kkaol59636jnfx
hostname(config)# crypto dynamic-map dyn1 1 set transform-set FirstSet
```

```
hostname(config)# crypto dynamic-map dyn1 1 set reverse-route
hostname(config)# crypto map mymap 1 ipsec-isakmp dynamic dyn1
hostname(config)# crypto map mymap interface outside
hostname(config)# write memory
```

Configuring Interfaces

A security appliance has at least two interfaces, referred to here as outside and inside. Typically, the outside interface is connected to the public Internet, while the inside interface is connected to a private network and is protected from public access.

To begin, configure and enable two interfaces on the security appliance. Then assign a name, IP address and subnet mask. Optionally, configure its security level, speed and duplex operation on the security appliance.

To configure interfaces, perform the following steps, using the command syntax in the examples:

-
- Step 1** To enter Interface configuration mode, in global configuration mode enter the **interface** command with the default name of the interface to configure. In the following example the interface is ethernet0.

```
hostname(config)# interface ethernet0
hostname(config-if)#
```

- Step 2** To set the IP address and subnet mask for the interface, enter the **ip address** command. In the following example the IP address is 10.10.4.100 and the subnet mask is 255.255.0.0.

```
hostname(config-if)# ip address 10.10.4.200 255.255.0.0
hostname(config-if)#
```

- Step 3** To name the interface, enter the **nameif** command, maximum of 48 characters. You cannot change this name after you set it. In the following example the name of the ethernet0 interface is outside.

```
hostname(config-if)# nameif outside
hostname(config-if)##
```

- Step 4** To enable the interface, enter the **no** version of the **shutdown** command. By default, interfaces are disabled.

```
hostname(config-if)# no shutdown
hostname(config-if)#
```

- Step 5** To save your changes, enter the **write memory** command.

```
hostname(config-if)# write memory
hostname(config-if)#
```

- Step 6** To configure a second interface, use the same procedure.
-

Configuring ISAKMP Policy and Enabling ISAKMP on the Outside Interface

The Internet Security Association and Key Management Protocol, also called IKE, is the negotiation protocol that lets two hosts agree on how to build an IPSec Security Association. Each ISAKMP negotiation is divided into two sections called Phase1 and Phase2.

Phase 1 creates the first tunnel to protect later ISAKMP negotiation messages. Phase 2 creates the tunnel that protects data travelling across the secure connection.

To set the terms of the ISAKMP negotiations, you create an ISAKMP policy. It includes the following:

- An authentication method, to ensure the identity of the peers.
- An encryption method, to protect the data and ensure privacy.
- A Hashed Message Authentication Codes method to ensure the identity of the sender and to ensure that the message has not been modified in transit.
- A Diffie-Hellman group to set the size of the encryption key.
- A time limit for how long the security appliance uses an encryption key before replacing it.

See [on page 27-3](#) in the “Configuring IPSec and ISAKMP” chapter of this guide for detailed information about the IKE policy keywords and their values.

To configure ISAKMP policies, in global configuration mode, enter the **isakmp policy** command with its various arguments. The syntax for ISAKMP policy commands is **isakmp policy priority attribute_name [attribute_value | integer]**.

Perform the following steps and use the command syntax in the following examples as a guide.

Step 1 Set the authentication method. The following example configures preshared key. The priority is 1 in this and all following steps.

```
hostname(config)# isakmp policy 1 authentication pre-share
hostname(config)#
```

Step 2 Set the encryption method. The following example configures 3DES.

```
hostname(config)# isakmp policy 1 encryption 3des
hostname(config)#
```

Step 3 Set the HMAC method. The following example configures SHA-1.

```
hostname(config)# isakmp policy 1 hash sha
hostname(config)#
```

Step 4 Set the Diffie-Hellman group. The following example configures Group 2.

```
hostname(config)# isakmp policy 1 group 2
hostname(config)#
```

Step 5 Set the encryption key lifetime. The following example configures 43,200 seconds (12 hours).

```
hostname(config)# isakmp policy 1 lifetime 43200
hostname(config)#
```

Step 6 Enable ISAKMP on the interface named outside.

```
hostname(config)# isakmp enable outside
hostname(config)#
```

- Step 7** To save your changes, enter the **write memory** command.

```
hostname(config)# write memory
hostname(config)#
```

Configuring an Address Pool

The security appliance requires a method for assigning IP addresses to users. A common method is using address pools. The alternatives are having a DHCP server assign address or having an AAA server assign them. The following example uses an address pool.

- Step 1** To configure an address pool, enter the **ip local pool** command. The syntax is **ip local pool poolname first_address-last_address**. In the following example the pool name is testpool.

```
hostname(config)# ip local pool testpool 192.168.0.10-192.168.0.15
hostname(config)#
```

- Step 2** Save your changes.

```
hostname(config)# write memory
hostname(config)#
```

Adding a User

To identify remote access users to the security appliance, configure usernames and passwords.

- Step 1** To add users, enter the **username** command. The syntax is **username username password password**. In the following example the username is testuser and the password is 12345678.

```
hostname(config)# username testuser password 12345678
hostname(config)#
```

- Step 2** Repeat Step 1 for each additional user.
-

Creating a Transform Set

A transform set combines an encryption method and an authentication method. During the IPsec security association negotiation with ISAKMP, the peers agree to use a particular transform set to protect a particular data flow. The transform set must be the same for both peers.

A transform set protects the data flows for the access list specified in the associated crypto map entry. You can create transform sets in the security appliance configuration, and then specify a maximum of 11 of them in a crypto map or dynamic crypto map entry. For more overview information, including a table that lists valid encryption and authentication methods, see [Creating a Transform Set](#) in [Chapter 36](#), “Configuring LAN-to-LAN IPsec VPNs” of this guide.

- Step 1** To configure a transform set, in global configuration mode enter the **crypto ipsec transform-set** command. The syntax is:

```
crypto ipsec transform-set transform-set-name encryption-method authentication-method
```

The following example configures a transform set with the name FirstSet, esp-3des encryption, and esp-md5-hmac authentication:

```
hostname(config)# crypto ipsec transform set FirstSet esp-3des esp-md5-hmac
hostname(config)#
```

- Step 2** Save the changes.

```
hostname(config)# write memory
hostname(config)#
```

Defining a Tunnel Group

A tunnel group is a set of records that contain tunnel connection policies. You configure a tunnel group to identify AAA servers, specify connection parameters, and define a default group policy. The security appliance stores tunnel groups internally.

There are two default tunnel groups in the security appliance system: DefaultRAGroup, which is the default IPSec remote-access tunnel group, and DefaultL2Lgroup, which is the default IPSec LAN-to-LAN tunnel group. You can change them but not delete them. The security appliance uses these groups to configure default tunnel parameters for remote access and LAN-to-LAN tunnel groups when there is no specific tunnel group identified during tunnel negotiation.

To establish a basic remote access connection, you must set three attributes for a tunnel group:

- Set the connection type to IPSec remote access.
- Configure the address assignment method, in the following example, address pool.
- Configure an authentication method, in the following example, preshared key.

- Step 1** To set the connection type to IPSec remote access, enter the **tunnel-group** command. The command syntax is **tunnel-group name type type**, where *name* is the name you assign to the tunnel group, and *type* is the type of tunnel. The tunnel types as you enter them in the CLI include the following:

- ipsec-ra (IPSec remote access)
- ipsec-l2l (IPSec LAN to LAN)

In the following example the name of the tunnel group is testgroup.

```
hostname(config)# tunnel-group testgroup type ipsec-ra
hostname(config)#
```

- Step 2** To configure an authentication method for the tunnel group, enter the general-attributes mode and then enter the **address-pool** command to create the address pool. In the following example the name of the group is testgroup and the name of the address pool is testpool.

```
hostname(config)# tunnel-group testgroup general-attributes
hostname(config-general)# address-pool testpool
```

- Step 3** To configure the authentication method, enter the ipsec-attributes mode and then enter the **pre-shared-key** command to create the preshared key. You need to use the same preshared key on both the security appliance and the client.

**Note**

The preshared key must be no larger than that used by the VPN client. If a Cisco VPN Client with a different preshared key size tries to connect to a security appliance, the client logs an error message indicating it failed to authenticate the peer.

The key is an alphanumeric string of 1-128 characters. In the following example the preshared key is 44kkaol59636jnfx.

```
hostname(config)# tunnel-group testgroup ipsec-attributes
hostname(config-ipsec)# pre-shared-key 44kkaol59636jnfx
```

- Step 4** Save your changes.

```
hostname(config)# write memory
hostname(config)#
```

Creating a Dynamic Crypto Map

The security appliance uses dynamic crypto maps to define a policy template where all the parameters do not have to be configured. These dynamic crypto maps let the security appliance receive connections from peers that have unknown IP addresses. Remote access clients fall in this category.

Dynamic crypto map entries identify the transform set for the connection. You also enable reverse routing, which lets the security appliance learn routing information for connected clients, and advertise it via RIP or OSPF.

- Step 1** To specify a transform set for a dynamic crypto map entry, enter the **crypto dynamic-map set transform-set** command.

The syntax is **crypto dynamic -map *dynamic-map-name seq-num set transform-set transform-set-name***. In the following example the name of the dynamic map is dyn1, the sequence number is 1, and the transform set name is FirstSet.

```
hostname(config)# crypto dynamic-map dyn1 1 set transform-set FirstSet
hostname(config)#
```

- Step 2** To enable RRI for any connection based on this crypto map entry, enter the **crypto dynamic-map set reverse route** command.

```
hostname(config)# crypto dynamic-map dyn1 1 set reverse-route
hostname(config)#
```

- Step 3** Save your changes.

```
hostname(config)# write memory
hostname(config)#
```

Creating a Crypto Map Entry to Use the Dynamic Crypto Map

Next create a crypto map entry that lets the security appliance use the dynamic crypto map to set the parameters of IPsec security associations.

In the following examples for this command, the name of the crypto map is mymap, the sequence number is 1, and the name of the dynamic crypto map is dyn1, which you created in the previous section, “[Creating a Dynamic Crypto Map](#).” Enter these commands in global configuration mode.

-
- Step 1** To create a crypto map entry that uses a dynamic crypto map, enter the **crypto map** command. The syntax is **crypto map** *map-name* *seq-num* **ipsec-isakmp** **dynamic** *dynamic-map-name*.

```
hostname(config)# crypto map mymap 1 ipsec-isakmp dynamic dyn1
hostname(config)#
```

- Step 2** To apply the crypto map to the outside interface, enter the **crypto map interface** command.

The syntax is **crypto map** *map-name* **interface** *interface-name*

```
hostname(config)# crypto map mymap interface outside
hostname(config)#
```



CHAPTER 33

Configuring Network Admission Control

This chapter includes the following sections.

- [Uses, Requirements, and Limitations, page 33-1](#)
- [Configuring Basic Settings, page 33-1](#)
- [Changing Advanced Settings, page 33-5](#)

Uses, Requirements, and Limitations

Network Admission Control (NAC) protects the enterprise network from intrusion and infection from worms, viruses, and rogue applications by performing endpoint compliancy and vulnerability checks as a condition for production access to the network. We refer to these checks as *posture validation*. You can configure posture validation to ensure that the anti-virus files, personal firewall rules, or intrusion protection software on a host establishing an IPSec session are up-to-date. Posture validation can include the verification that the applications running on the remote hosts are updated with the latest patches. NAC supplements the identity-based validation that IPSec and other access methods provide. It is especially useful for protecting the enterprise network from hosts that are not subject to automatic network policy enforcement, such as home PCs.



Note

When configured to support NAC, the security appliance functions as a client of a Cisco Secure Access Control Server, requiring that you install a minimum of one Access Control Server on the network to provide NAC authentication services.

Following the configuration of one or more Access Control Servers on the network, you must use the **aaa-server** command to name the Access Control Server group. Then follow the instructions in [Configuring Basic Settings, page 33-1](#) to configure NAC.

ASA support for NAC is limited to remote access IPSec and L2TP over IPSec sessions. NAC on the ASA does not support WebVPN, non-VPN traffic, IPv6, and multimode.

Configuring Basic Settings

The instructions in the following sections describe how to enter the minimum set of commands to configure support for NAC on the security appliance:

- [Specifying the Access Control Server Group, page 33-2](#)

- [Enabling NAC, page 33-2](#)
- [Configuring the Default ACL for NAC, page 33-3](#)
- [Configuring Exemptions from NAC, page 33-4](#)

**Note**

See [Uses, Requirements, and Limitations, page 33-1](#) before following these instructions.

Specifying the Access Control Server Group

You must configure at least one Cisco Access Control Server to support NAC. Then use the **aaa-server host** command to name the Access Control Server group even if the group contains only one server. Then enter the following command in tunnel-group general-attributes configuration mode to specify the same group as the group to be used for NAC posture validation:

```
nac-authentication-server-group server-group
```

server-group must match the server-tag variable specified in the **aaa-server host** command.

For example, enter the following command to specify acs-group1 as the authentication server group to be used for NAC posture validation:

```
hostname(config-group-policy) # nac-authentication-server-group acs-group1  
hostname(config-group-policy)
```

To inherit the authentication server group from the default remote access group, access the alternative group policy from which to inherit it, then enter the following command:

```
no nac-authentication-server-group
```

For example:

```
hostname(config-group-policy) # no nac-authentication-server-group  
hostname(config-group-policy)
```

Enabling NAC

To enable or disable NAC for a group policy, enter the following command in group-policy configuration mode:

```
nac {enable | disable}
```

The following example enables NAC for the group policy:

```
hostname(config-group-policy) # nac enable  
hostname(config-group-policy)
```

To inherit the NAC setting from the default group policy, access the alternative group policy from which to inherit it, then issue the following command:

```
no nac
```

For example:

```
hostname(config-group-policy) # no nac
```

```
hostname(config-group-policy) #
```

Configuring the Default ACL for NAC

Each group policy points to a default ACL to be applied to hosts that match the policy and are eligible for NAC. The security appliance applies the NAC default ACL before posture validation. Following posture validation, the security appliance replaces the default ACL with the one obtained from the Access Control Server for the remote host. It retains the default ACL if posture validation fails.

The security appliance also applies the NAC default ACL if clientless authentication is enabled (which is the default setting).



Note

Because NAC is disabled by default, VPN traffic traversing the security appliance is not subject to the NAC Default ACL until NAC is enabled.

Enter the following command in group-policy configuration mode to specify the ACL to be used as the default ACL for NAC sessions:

```
nac-default-acl value acl-name
```

acl-name is the name of the posture validation server group, as configured on the security appliance using the **aaa-server host** command. The name must match the server-tag variable specified in that command.

For example, enter the following command to specify *acl-1* as the NAC default ACL:

```
hostname(config-group-policy) # nac-default-acl value acl-1  
hostname(config-group-policy)
```

To inherit the ACL from the default group policy, access the alternative group policy from which to inherit it and enter the following command.

```
no nac-default-acl
```

For example:

```
hostname(config-group-policy) # no nac-default-acl  
hostname(config-group-policy)
```

You also have the option of disinherit the ACL from the default group policy and specifying no NAC default ACL. To do so, enter the following command:

```
nac-default-acl none
```

For example:

```
hostname(config-group-policy) # nac-default-acl none  
hostname(config-group-policy)
```

Configuring Exemptions from NAC

The security appliance configuration stores a list of exemptions from NAC posture validation. You can specify the operating systems that are exempt. If you specify an ACL, the client running the operating system specified is exempt from posture validation and the client traffic is subject to the ACL.

To add an entry to the list of remote computer types that are exempt from NAC posture validation, enter the following command in group-policy configuration mode:

```
vpn-nac-exempt os "os name" [filter acl-name] [disable]
```

**Note**

This command does not overwrite the previously added entry to the exception list; enter the command once for each operating system and ACL you want to exempt.

os name is the operating system name. Use quotation marks if the name includes a space (for example, "Windows XP").

For example, enter the following command to add all hosts running Windows XP to the list of computers that are exempt from posture validation:

```
hostname(config-group-policy)# vpn-nac-exempt os "Windows XP"  
hostname(config-group-policy)
```

The remaining keywords and arguments are optional:

- **filter** to apply an ACL to filter the traffic if the computer matches the os name.
- *acl-name* is the name of the ACL present in the security appliance configuration.
- **disable** to disable the entry in the exemption list without removing it from the list. Not entering this keyword enables the entry.

For example, enter the following command to exempt all hosts running Windows 98 and apply the ACL *acl-1* to traffic from those hosts:

```
hostname(config-group-policy)# vpn-nac-exempt os "Windows 98" filter acl-1  
hostname(config-group-policy)
```

The following example shows how to add the same entry to the exemption list, but disable it:

```
hostname(config-group-policy)# vpn-nac-exempt os "Windows 98" filter acl-1 disable  
hostname(config-group-policy)
```

To disable inheritance and specify that all hosts are subject to posture validation, enter the following command:

```
vpn-nac-exempt none
```

For example:

```
hostname(config-group-policy)# no vpn-nac-exempt none  
hostname(config-group-policy)
```

To remove an entry from the exemption list, enter the following command, naming the operating system (and ACL) in the exemption to be removed.

```
no vpn-nac-exempt [os "os name"] [filter acl-name]
```

For example, enter the following command to remove the entry with Windows 98 and acl-1 from the exemption list, regardless of whether it is disabled:

```
hostname(config-group-policy) # no vpn-nac-exempt os "Windows 98" filter acl-1
hostname(config-group-policy)
```

To remove all entries from the exemption list associated with this group policy and inherit the list from the default group policy, enter the following command without specifying additional keywords:

```
no vpn-nac-exempt
```

For example:

```
hostname(config-group-policy) # no vpn-nac-exempt
hostname(config-group-policy)
```

Changing Advanced Settings

The security appliance provides default settings for NAC. Use the instructions in this section to adjust these settings for adherence to the policies in force in your network.

Changing Clientless Authentication Settings

NAC support for clientless authentication is configurable. It applies to hosts that do not have a posture agent, such as the Cisco Trust Agent. The security appliance applies the default access policy, sends the EAP over UDP request for posture validation, and the request times out. If the security appliance is not configured to request a policy for clientless hosts from the Access Control Server, it retains the default access policy already in use for the clientless host. If the security appliance is configured to request a policy for clientless hosts from the Access Control Server, it does so and the Access Control Server downloads the access policy to be enforced by the security appliance.

Enabling and Disabling Clientless Authentication

Enter the following command in global configuration mode to enable clientless authentication:

```
euo allow clientless
```

For example:

```
hostname(config) # euo allow clientless
hostname(config) #
```

The **euo clientless** command is meaningful only if NAC is enabled.



Note

Clientless authentication is enabled by default.

Enter the following command in global configuration mode to disable clientless authentication:

```
no euo allow clientless
```

For example:

```
hostname(config)# no eou allow clientless
hostname(config)#
```

Changing the Login Credentials Used for Clientless Authentication

When clientless authentication is enabled, and the security appliance fails to receive a response to a validation request from the remote host, it sends a clientless authentication request on behalf of the remote host to the Access Control Server. The request includes the login credentials that match those configured for clientless authentication on the Access Control Server. The default username and password for clientless authentication on the security appliance matches the default username and password on the Access Control Server; the default username and password are both “clientless”. If you change these values on the Access Control Server, you must also do so on the security appliance.

Enter the following command in global configuration mode to change the username used for clientless authentication:

eou clientless username *username*

username must match the username configured on the Access Control Server to support clientless hosts. Enter 1 to 64 ASCII characters, excluding leading and trailing spaces, pound signs (#), question marks (?), quotation marks ("), asterisks (*), and angle brackets (< and >).

Enter the following command in global configuration mode to change the password used for clientless authentication:

eou clientless password *password*

password must match the password configured on the Access Control Server to support clientless hosts. Enter 4 – 32 ASCII characters.

You can specify only the username, only the password, or both. For example, enter the following commands to change the username and password for clientless authentication to sherlock and 221B-baker, respectively:

```
hostname(config)# eou clientless username sherlock
hostname(config)# eou clientless password 221B-baker
hostname(config)#
```

To change the username to its default value, enter the following command:

no eou clientless username

For example:

```
hostname(config)# no eou clientless username
hostname(config)#
```

To change the password to its default value, enter the following command:

no eou clientless password

For example:

```
hostname(config)# no eou clientless password
hostname(config)#
```

Configuring NAC Session Attributes

The ASA provides default settings for the attributes that specify communications between the security appliance and the remote host. These attributes specify the port no. to communicate with posture agents on remote hosts and the expiration counters that impose limits on the communications with the posture agents. These attributes, the default settings, and the commands you can enter to change them are as follows:

- Port no. on the client endpoint to be used for EAP over UDP communication with posture agents.

The default port no. is 21862. Enter the following command in global communication mode to change it:

eou port *port_number*

port_number must match the port number configured on the CTA. Enter a value in the range 1024 to 65535.

For example, enter the following command to change the port number for EAP over UDP communication to 62445:

```
hostname(config)# eou port 62445  
hostname(config)#
```

To change the port number to its default value, use the **no** form of this command, as follows:

no eou port

For example:

```
hostname(config)# no eou port  
hostname(config)#
```

- Retransmission retry timer

When the security appliance sends an EAP over UDP message to the remote host, it waits for a response. If it fails to receive a response within *n* seconds, it resends the EAP over UDP message. By default, the retransmission timer is 3 seconds. To change this value, enter the following command in global configuration mode:

eou timeout retransmit *seconds*

seconds is a value in the range 1 to 60.

The following example changes the retransmission timer to 6 seconds:

```
hostname(config)# eou timeout retransmit 6  
hostname(config)#
```

To change the retransmission retry timer to its default value, use the **no** form of this command, as follows:

no eou timeout retransmit

For example:

```
hostname(config)# no eou timeout retransmit  
hostname(config)#
```

- Retransmission retries

When the security appliance sends an EAP over UDP message to the remote host, it waits for a response. If it fails to receive a response, it resends the EAP over UDP message. By default, it retries up to 3 times. To change this value, enter the following command in global configuration mode:

euo max-retry *retries*

retries is a value in the range 1 to 3.

The following example limits the number of EAP over UDP retransmissions to 1:

```
hostname(config)# euo max-retry 1
hostname(config)#
```

To change the maximum number of retransmission retries to its default value, use the **no** form of this command, as follows:

no euo max-retry

For example:

```
hostname(config)# no euo max-retry
hostname(config)#
```

- Session reinitialization timer

When the retransmission retry counter matches the max-retry value, the security appliance terminates the EAP over UDP session with the remote host and starts the hold timer. When the hold timer equals *n* seconds, the security appliance establishes a new EAP over UDP session with the remote host. By default, the maximum number of seconds to wait before establishing a new session is 180 seconds. To change this value, enter the following command in global configuration mode:

euo timeout hold-period *seconds*

seconds is a value in the range 60 to 86400.

For example, enter the following command to change the wait period before initiating a new EAP over UDP association to 120 seconds:

```
hostname(config)# euo timeout hold-period 120
hostname(config)#
```

To change the session reinitialization to its default value, use the **no** form of this command, as follows:

no euo timeout hold-period

For example:

```
hostname(config)# no euo timeout hold-period
hostname(config)#
```

Setting the Query-for-Posture-Changes Timer

After each successful posture validation, the security appliance starts a status query timer. The expiration of this timer triggers a query to the remote host for changes in posture since the last posture validation. A response indicating no change resets the status query timer. A response indicating a change in posture triggers an unconditional posture revalidation. The security appliance maintains the current access policy during revalidation.

By default, the interval between each successful posture validation and the status query, and each subsequent status query, is 300 seconds (5 minutes). The group policy inherits the value of the status query timer from the default group policy unless you change it. Enter the following command in group-policy configuration mode to change the status query interval:

nac-sq-period *seconds*

seconds must be in the range is 300 to 1800 seconds (5 to 30 minutes).

The following example changes the status query timer to 1800 seconds:

```
hostname(config-group-policy)# nac-sq-period 1800
hostname(config-group-policy)
```

To inherit the value of the status query timer from the default group policy, access the alternative group policy from which to inherit it, then enter the following command.

no nac-sq-period [*seconds*]

For example:

```
hostname(config-group-policy)# no nac-sq-period
hostname(config-group-policy)
```

Setting the Revalidation Timer

After each successful posture validation, the security appliance starts a revalidation timer. The expiration of this timer triggers the next unconditional posture validation. The security appliance maintains the current access policy during revalidation.

By default, the interval between each successful posture validation is 36000 seconds (10 hours). The group policy inherits the value of the revalidation timer from the default group policy unless you change it. Enter the following command in group-policy configuration mode to change the revalidation interval:

nac-reval-period *seconds*

seconds must be in the range is 300 to 86400 seconds (5 minutes to 24 hours).

For example, enter the following command to change the revalidation timer to 86400 seconds:

```
hostname(config-group-policy)# nac-reval-period 86400
hostname(config-group-policy)
```

To inherit the value of the revalidation timer from the default group policy, access the alternative group policy from which to inherit it, then enter the following command.

no nac-reval-period

For example:

```
hostname(config-group-policy)# no nac-reval-period
hostname(config-group-policy)
```




CHAPTER 34

Configuring Easy VPN Services on the ASA 5505

This chapter describes how to configure the ASA 5505 as an Easy VPN hardware client. This chapter assumes you have configured the switch ports and VLAN interfaces of the ASA 5505 (see [Chapter 4, “Configuring Switch Ports and VLAN Interfaces for the Cisco ASA 5505 Adaptive Security Appliance”](#)).



Note

The Easy VPN hardware client configuration specifies the IP address of its primary and secondary (backup) Easy VPN servers. Any ASA, including another ASA 5505 configured as a headend, a VPN 3000 Series Concentrator, an IOS-based router, or a firewall can act as an Easy VPN server. An ASA 5505 cannot, however function as both a client and a server simultaneously. To configure an ASA 5505 as a server, see [“Specifying the Client/Server Role of the Cisco ASA 5505” section on page 34-1](#). Then configure the ASA 5505 as you would any other ASA, beginning with the [“Getting Started” section on page 2-1](#) of this guide.

This chapter includes the following sections:

- [Specifying the Client/Server Role of the Cisco ASA 5505, page 34-1](#)
- [Specifying the Primary and Secondary Servers, page 34-2](#)
- [Specifying the Mode, page 34-3](#)
- [Configuring Automatic Xauth Authentication, page 34-4](#)
- [Configuring IPSec Over TCP, page 34-4](#)
- [Comparing Tunneling Options, page 34-5](#)
- [Specifying the Tunnel Group or Trustpoint, page 34-6](#)
- [Configuring Split Tunneling, page 34-7](#)
- [Configuring Device Pass-Through, page 34-8](#)
- [Configuring Remote Management, page 34-8](#)
- [Guidelines for Configuring the Easy VPN Server, page 34-9](#)

Specifying the Client/Server Role of the Cisco ASA 5505

The Cisco ASA 5505 can function as a Cisco Easy VPN hardware client (also called “Easy VPN Remote”) or as a server (also called a “headend”), but not both at the same time. It does not have a default role. Use one of the following commands in global configuration mode to specify its role:

- `vpnclient enable` to specify the role of the ASA 5505 as an Easy VPN Remote
- **no `vpnclient enable`** to specify the role of the ASA 5505 as server

The following example shows how to specify the ASA 5505 as an Easy VPN hardware client:

```
hostname(config)# vpnclient enable
hostname(config)#
```

The CLI responds with an error message indicating that you must remove certain data elements if you switch from server to hardware client, depending on whether the elements are present in the configuration. [Table 34-1](#) lists the data elements that are permitted in both client and server configurations, and not permitted in client configurations.

Table 34-1 Configuration Privileges and Restrictions on the ASA 5505

Permitted in Both Client and Server Configurations	Not Permitted in Client Configurations
crypto ca trustpoints	tunnel-groups
digital certificates	isakmp policies
group-policies	crypto maps
crypto dynamic-maps	
crypto ipsec transform-sets	
crypto ipsec security-association lifetime	
crypto ipsec fragmentation before-encryption	
crypto ipsec df-bit copy-df	

An ASA 5505 configured as an Easy VPN hardware client retains the commands listed in the first column within its configuration, however, some have no function in the client role.

The following example shows how to specify the ASA 5505 as an Easy VPN server:

```
hostname(config)# no vpnclient enable
hostname(config)#
```

After entering the no version of this command, configure the ASA 5505 as you would any other ASA, beginning with [“Getting Started” section on page 2-1](#) of this guide.

Specifying the Primary and Secondary Servers

Before establishing a connection with an Easy VPN hardware client, you must specify the IP address of an Easy VPN server to which it will connect. Any ASA, including another ASA 5505 configured as a headend, a VPN 3000 Series Concentrator, an IOS-based router, or a firewall can act as an Easy VPN server. Use the **vpnclient server** command in global configuration mode, as follows:

[no] vpnclient server *ip_primary* [*ip_secondary_1...ip_secondary_10*]

no removes the command from the running configuration.

ip_primary_address is the IP address or DNS name of the primary Easy VPN server.

ip_secondary_address_n (Optional) is a list of the IP addresses or DNS names of up to ten backup Easy VPN servers. Use a space to separate the items in the list.

For example, enter the following command to configure a VPN client to use Easy VPN Server 10.10.10.15 as the primary server, and 10.10.10.30 and 192.168.10.45 as alternate servers:

```
hostname(config)# vpnclient server 10.10.10.15 10.10.10.30 192.168.10.10
hostname(config)#
```

Specifying the Mode

The Easy VPN Client supports one of two modes of operation: Client Mode or Network Extension Mode (NEM). The mode of operation determines whether the inside hosts relative to the Easy VPN Client are accessible from the Enterprise network over the tunnel. Specifying a mode of operation is mandatory before making a connection because Easy VPN Client does not have a default mode.

Client mode, also called Port Address Translation (PAT) mode, isolates the IP addresses of all devices on the Easy VPN Client private network from those on the enterprise network. The Easy VPN Client performs PAT for all VPN traffic for its inside hosts. IP address management is neither required for the Easy VPN Client inside interface or the inside hosts.

NEM makes the inside interface and all inside hosts routeable across the enterprise network over the tunnel. Hosts on the inside network obtain their IP addresses from an accessible subnet (statically or via DHCP) pre-configured with static IP addresses. PAT does not apply to VPN traffic in NEM. This mode does not require a VPN configuration for each client. The Cisco ASA 5505 configured for NEM mode supports automatic tunnel initiation. The configuration must store the group name, user name, and password. Automatic tunnel initiation is disabled if secure unit authentication is enabled.



Note

If the Easy VPN hardware client is using NEM and has connections to secondary servers, use the **crypto map set reverse-route** command on each headend device to configure dynamic announcements of the remote network using Reverse Route Injection (RRI).

To specify the mode for Easy VPN Clients, enter the following command in configuration mode:

```
[no] vpnclient mode {client-mode | network-extension-mode}
```

no removes the command from the running configuration.

NEM with Multiple Interfaces

If you have an ASA 5505 security appliance (version 7.2 (3) or higher) configured as an Easy VPN Client in Network Extension Mode with multiple interfaces configured, the security appliance builds a tunnel for locally encrypted traffic only from the interface with the highest security level.

For example, consider the following configuration:

```
vlan1 security level 100 nameif inside
vlan2 security level 0 nameif outside
vlan12 security level 75 nameif work
```

In this scenario, the security appliance builds the tunnel only for vlan1, the interface with the highest security level. If you want to encrypt traffic from vlan12, you must change the security level of interface vlan1 to a lower value than that of vlan 12.

Configuring Automatic Xauth Authentication

The ASA 5505 configured as an Easy VPN hardware client automatically authenticates when it connects to the Easy VPN server if all of the following conditions are true:

- Secure unit authentication is disabled on the server.
- The server requests IKE Extended Authenticate (Xauth) credentials.

Xauth provides the capability of authenticating a user within IKE using TACACS+ or RADIUS. Xauth authenticates a user (in this case, the Easy VPN hardware client) using RADIUS or any of the other supported user authentication protocols.

- The client configuration contains an Xauth username and password.

Enter the following command in global configuration mode to configure the Xauth username and password:

```
vpnclient username xauth_username password xauth_password
```

You can use up to 64 characters for each.

For example, enter the following command to configure the Easy VPN hardware client to use the XAUTH username testuser and password ppurkml:

```
hostname(config)# vpnclient username testuser password ppurkml
hostname(config)#
```

To remove the username and password from the running configuration, enter the following command:

```
no vpnclient username
```

For example:

```
hostname(config)# no vpnclient username
hostname(config)#
```

Configuring IPSec Over TCP

By default, the Easy VPN hardware client and server encapsulate IPSec in User Datagram Protocol (UDP) packets. Some environments, such as those with certain firewall rules, or NAT and PAT devices, prohibit UDP. To use standard Encapsulating Security Protocol (ESP, Protocol 50) or Internet Key Exchange (IKE, UDP 500) in such environments, you must configure the client and the server to encapsulate IPSec within TCP packets to enable secure tunneling. If your environment allows UDP, however, configuring IPSec over TCP adds unnecessary overhead.

To configure the Easy VPN hardware client to use TCP-encapsulated IPSec, enter the following command in global configuration mode:

```
vpnclient ipsec-over-tcp [port tcp_port]
```

The Easy VPN hardware client uses port 10000 if the command does not specify a port number.

If you configure an ASA 5505 to use TCP-encapsulated IPSec, enter the following command to let it send large packets over the outside interface:

```
hostname(config)# crypto ipsec df-bit clear-df outside
hostname(config)#
```

This command clears the Don't Fragment (DF) bit from the encapsulated header. A DF bit is a bit within the IP header that determines whether the packet can be fragmented. This command lets the Easy VPN hardware client send packets that are larger than the MTU size.

The following example shows how to configure the Easy VPN hardware client to use TCP-encapsulated IPSec, using the default port 10000, and to let it send large packets over the outside interface:

```
hostname(config)# vpnclient ipsec-over-tcp
hostname(config)# crypto ipsec df-bit clear-df outside
hostname(config)#
```

The next example shows how to configure the Easy VPN hardware client to use TCP-encapsulated IPSec, using the port 10501, and to let it send large packets over the outside interface:

```
hostname(config)# vpnclient ipsec-over-tcp port 10501
hostname(config)# crypto ipsec df-bit clear-df outside
hostname(config)#
```

To remove the attribute from the running configuration, use the **no** form of this command, as follows:

no vpnclient ipsec-over-tcp

For example:

```
hostname(config)# no vpnclient ipsec-over-tcp
hostname(config)#
```

Comparing Tunneling Options

The tunnel types the Cisco ASA 5505 configured as an Easy VPN hardware client sets up depends on a combination of the following factors:

- Use of the **split-tunnel-network-list** and the **split-tunnel-policy** commands on the headend to permit, restrict, or prohibit split tunneling. (See the [Creating a Network List for Split-Tunneling](#), page 30-42 and “Setting the Split-Tunneling Policy” section on page 30-42, respectively.)

Split tunneling determines the networks for which the remote-access client encrypts and sends data through the secured VPN tunnel, and determines which traffic it sends to the Internet in the clear.

- Use of the **vpnclient management** command to specify one of the following automatic tunnel initiation options:
 - **tunnel** to limit administrative access to the client side by specific hosts or networks on the corporate side and use IPSec to add a layer of encryption to the management sessions over the HTTPS or SSH encryption that is already present.
 - **clear** to permit administrative access using the HTTPS or SSH encryption used by the management session.
 - **no** to prohibit management access



Caution


Cisco does not support the use of the **vpnclient management** command if a NAT device is present between the client and the Internet.

- Use of the **vpnclient mode** command to specify one of the following modes of operation:
 - **client** to use Port Address Translation (PAT) mode to isolate the addresses of the inside hosts, relative to the client, from the enterprise network.

- **network-extension-mode** to make those addresses accessible from the enterprise network.

Figure 34-1 shows the types of tunnels that the Easy VPN client initiates, based on the combination of the commands you enter.

Figure 34-1 Easy VPN Hardware Client Tunneling Options for the Cisco ASA 5505



Configuration factors:

1. Certs or Preshare Keys (Phase 1- main mode or aggressive mode)
2. Mode: Client or NEM
3. All-or-nothing or Split-tunneling
4. Management Tunnels
5. IUA to VPN3000 or ASA headend

* Only for ASA or VPN3000 Headends

153780

The term “All-Or-Nothing” refers to the presence or absence of an access list for split tunneling. The access list (“ST-list”) distinguishes networks that require tunneling from those that do not.

Specifying the Tunnel Group or Trustpoint

When configuring the Cisco ASA 5505 as an Easy VPN hardware client, you can specify a tunnel group or trustpoint configured on the Easy VPN server, depending on the Easy VPN server configuration. See the section that names the option you want to use:

- [Specifying the Tunnel Group](#)
- [Specifying the Trustpoint](#)

Specifying the Tunnel Group

Enter the following command in global configuration mode to specify the name of the VPN tunnel group and password for the Easy VPN client connection to the server:

```
vpnclient vpngroup group_name password preshared_key
```

group_name is the name of the VPN tunnel group configured on the Easy VPN server. You must configure this tunnel group on the server before establishing a connection.

preshared_key is the IKE pre-shared key used for authentication on the Easy VPN server.

For example, enter the following command to identify the VPN tunnel group named TestGroup1 and the IKE preshared key my_key123.

```
hostname(config)# vpnclient vpngroup TestGroup1 password my_key123  
hostname(config)#
```

To remove the attribute from the running configuration, enter the following command:

```
no vpnclient vpngroup
```

If the configuration of the ASA 5505 running as an Easy VPN client does not specify a tunnel group, the client attempts to use an RSA certificate.

For example:

```
hostname(config)# no vpnclient vpngroup  
hostname(config)#
```

Specifying the Trustpoint

A trustpoint represents a CA identity, and possibly a device identity, based on a certificate the CA issues. These parameters specify how the security appliance obtains its certificate from the CA and define the authentication policies for user certificates issued by the CA.

First define the trustpoint using the **crypto ca trustpoint** command, as described in [“Configuring Trustpoints” section on page 39-7](#). Then enter the following command in global configuration mode to name the trustpoint identifying the RSA certificate to use for authentication:

```
vpnclient trustpoint trustpoint_name [chain]
```

trustpoint_name names the trustpoint identifying the RSA certificate to use for authentication.

(Optional) **chain** sends the entire certificate chain.

For example, enter the following command to specify the identity certificate named central and send the entire certificate chain:

```
hostname(config)# crypto ca trustpoint central  
hostname(config)# vpnclient trustpoint central chain  
hostname(config)#
```

To remove the attribute from the running configuration, enter the following command:

```
no vpnclient trustpoint
```

For example:

```
hostname(config)# no vpnclient trustpoint  
hostname(config)#
```

Configuring Split Tunneling

Split tunneling lets a remote-access IPSec client conditionally direct packets over an IPSec tunnel in encrypted form or to a network interface in clear text form.

The Easy VPN server pushes the split tunneling attributes from the group policy to the Easy VPN Client for use only in the work zone. See [Configuring Split-Tunneling Attributes, page 30-42](#) to configure split tunneling on the Cisco ASA 5505.

Enter the following command in global configuration mode to enable the automatic initiation of IPSec tunnels when NEM and split tunneling are configured:

```
[no] vpnclient nem-st-autoconnect
```

no removes the command from the running configuration.

For example:

```
hostname(config)# vpnclient nem-st-autoconnect
hostname(config)#
```

Configuring Device Pass-Through

Devices such as Cisco IP phones, wireless access points, and printers are incapable of performing authentication. Enter the following command in global configuration mode to exempt such devices from authentication, thereby providing network access to them, if individual user authentication is enabled:

```
[no] vpnclient mac-exempt mac_addr_1 mac_mask_1 [mac_addr_2 mac_mask_2...mac_addr_n
mac_mask_n]
```

no removes the command from the running configuration.

mac_addr is the MAC address, in dotted hexadecimal notation, of the device to bypass individual user authentication.

mac_mask is the network mask for the corresponding MAC address. A MAC mask of ffff.ff00.0000 matches all devices made by the same manufacturer. A MAC mask of ffff.ffff.ffff matches a single device.

Only the first six characters of the specific MAC address are required if you use the MAC mask ffff.ff00.0000 to specify all devices by the same manufacturer. For example, Cisco IP phones have the Manufacturer ID 00036b, so the following command exempts any Cisco IP phone, including Cisco IP phones, you might add in the future:

```
hostname(config)# vpnclient mac-exempt 0003.6b00.0000 ffff.ff00.0000
hostname(config)#
```

The next example provides greater security but less flexibility because it exempts one specific Cisco IP phone:

```
hostname(config)# vpnclient mac-exempt 0003.6b54.b213 ffff.ffff.ffff
hostname(config)#
```

Configuring Remote Management

The Cisco ASA 5505, operating as an Easy VPN hardware client, supports management access using SSH or HTTPS, with or without a second layer of additional encryption. You can configure the Cisco ASA 5505 to require IPSec encryption within the SSH or HTTPS encryption.

Use the **vpnclient management clear** command in global configuration mode to use normal routing to provide management access from the corporate network to the outside interface of the ASA 5505 (no tunneling management packets).

**Caution**

Do not configure a management tunnel on a Cisco ASA 5505 configured as an Easy VPN hardware client if a NAT device is operating between the Easy VPN hardware client and the Internet. In that configuration, use the **vpnclient management clear** command.

Use the **vpnclient management tunnel** command in global configuration mode if you want to automate the creation of IPSec tunnels to provide management access from the corporate network to the outside interface of the ASA 5505. The Easy VPN hardware client and server create the tunnels automatically after the execution of the **vpnclient server** command. The syntax of the **vpnclient management tunnel** command follows:

```
vpnclient management tunnel ip_addr_1 ip_mask_1 [ip_addr_2 ip_mask_2...ip_addr_n ip_mask_n]
```

For example, enter the following command to automate the creation of an IPSec tunnel to provide management access to the host with IP address 192.168.10.10:

```
hostname(config)# vpnclient management tunnel 192.198.10.10 255.255.255.0
hostname(config)#
```

The **no** form of this command sets up IPSec for management tunnels in accordance with the **split-tunnel-policy** and **split-tunnel-network-list** commands.

no vpnclient management

For example:

```
hostname(config)# no vpnclient management
hostname(config)#
```

Guidelines for Configuring the Easy VPN Server

The following sections address the Easy VPN hardware client considerations that apply to the Easy VPN server:

- [Group Policy and User Attributes Pushed to the Client](#)
- [Authentication Options](#)

Group Policy and User Attributes Pushed to the Client

Upon tunnel establishment, the Easy VPN server pushes the values of the group policy or user attributes stored in its configuration to the Easy VPN hardware client. Therefore, to change certain attributes pushed to the Easy VPN hardware client, you must modify them on the security appliances configured as the primary and secondary Easy VPN servers. This section identifies the group policy and user attributes pushed to the Easy VPN hardware client.

**Note**

This section serves only as a reference. For complete instructions on configuring group policies and users, see [Configuring Tunnel Groups, Group Policies, and Users, page 30-1](#).

Use [Table 34-2](#) as a guide for determining which commands to enter to modify the group policy or user attributes.

Table 34-2 *Group Policy and User Attributes Pushed to the Cisco ASA 5505 Configured as an EasyVPN Hardware Client*

Command	Description
backup-servers	Sets up backup servers on the client in case the primary server fails to respond.
banner	Sends a banner to the client after establishing a tunnel.
client-access-rule	Applies access rules.
client-firewall	Sets up the firewall parameters on the VPN client.
default-domain	Sends a domain name to the client.
dns-server	Specifies the IP address of the primary and secondary DNS servers, or prohibits the use of DNS servers.
dhcp-network-scope	Specifies the IP subnetwork to which the DHCP server assigns address to users within this group.
group-lock	Specifies a tunnel group to ensure that users connect to that group.
ipsec-udp	Uses UDP encapsulation for the IPSec tunnels.
ipsec-udp-port	Specifies the port number for IPSec over UDP.
nem	Enables or disables network extension mode.
password-storage	Lets the VPN user save a password in the user profile.
pfs	Commands the VPN client to use perfect forward secrecy.
re-xauth	Requires XAUTH authentication when IKE rekeys. Note: Disable re-xauth if secure unit authentication is enabled.
secure-unit-authentication	Enables interactive authentication for VPN hardware clients.
split-dns	Pushes a list of domains for name resolution.
split-tunnel-network-list	Specifies one of the following: - No access list exists for split tunneling. All traffic travels across the tunnel. - Identifies the access list the security appliance uses to distinguish networks that require tunneling and those that do not. Split tunneling lets a remote-access IPSec client conditionally direct packets over an IPSec tunnel in encrypted form, or to a network interface in cleartext form. With split-tunneling enabled, packets not bound for destinations on the other side of the IPSec tunnel do not have to be encrypted, sent across the tunnel, decrypted, and then routed to a final destination.

Table 34-2 *Group Policy and User Attributes Pushed to the Cisco ASA 5505 Configured as an EasyVPN Hardware Client (continued)*

Command	Description
split-tunnel-policy	Lets a remote-access IPSec client conditionally direct packets over an IPSec tunnel in encrypted form, or to a network interface in cleartext form. Options include the following: - split-tunnel-policy—Indicates that you are setting rules for tunneling traffic. - excludespecified—Defines a list of networks to which traffic goes in the clear. - tunnelall—Specifies that no traffic goes in the clear or to any other destination than the Easy VPN server. Remote users reach Internet networks through the corporate network and do not have access to local networks. - tunnelspecified—Tunnels all traffic from or to the specified networks. This option enables split tunneling. It lets you create a network list of addresses to tunnel. Data to all other addresses travels in the clear, and is routed by the remote user's internet service provider.
user-authentication	Enables individual user authentication for hardware-based VPN clients.
vpn-access-hours	Restricts VPN access hours.
vpn-filter	Applies a filter to VPN traffic.
vpn-idle-timeout	Specifies the number of minutes a session can be idle before it times out.
vpn-session-timeout	Specifies the maximum number of minutes for VPN connections.
vpn-simultaneous-logins	Specifies the maximum number of simultaneous logins.
vpn-tunnel-protocol	Specifies the permitted tunneling protocols.
wins-server	Specifies the IP address of the primary and secondary WINS servers, or prohibits the use of WINS servers.

**Note**

IPSec NAT-T connections are the only IPSec connection types supported on the home VLAN of a Cisco ASA 5505. IPSec over TCP and native IPSec connections are not supported.

Authentication Options

The ASA 5505 supports the following authentication mechanisms, which it obtains from the group policy stored on the Easy VPN Server. The following list identifies the authentication options supported by the Easy VPN hardware client, however, you must configure them on the Easy VPN server:

- Secure unit authentication (SUA, also called Interactive unit authentication)

Ignores the **vpnclient username** Xauth command (described in [“Configuring Automatic Xauth Authentication” section on page 34-4](#)) and requires the user to authenticate the ASA 5505 by entering a password. By default, SUA is disabled. You can use the **secure-unit-authentication enable** command in group-policy configuration mode to enable SUA. See [Configuring Secure Unit Authentication, page 30-45](#).

- Individual user authentication

Requires users behind the ASA 5505 to authenticate before granting them access to the enterprise VPN network. By default, IUA is disabled. To enable IUA, Use the **user-authentication enable** command in group-policy configuration mode. See [Configuring User Authentication, page 30-45](#).

The security appliance works correctly from behind a NAT device, and if the ASA5505 is configured in NAT mode, the provisioned IP (to which the clients all PAT) is injected into the routing table on the central-site device.

**Caution**

Do not configure IUA on a Cisco ASA 5505 configured as an Easy VPN server if a NAT device is operating between the server and the Easy VPN hardware client.

To set or remove the idle timeout period after which the Easy VPN Server terminates the client's access, use the **user-authentication-idle-timeout** command. See [Configuring an Idle Timeout, page 30-46](#).

- Authentication by HTTP redirection

The Cisco Easy VPN server intercepts HTTP traffic and redirects the user to a login page if one of the following is true:

- SUA or the username and password are not configured on the Easy VPN hardware client.
- IAU is enabled.

HTTP redirection is automatic and does not require configuration on the Easy VPN Server.

- Preshared keys, digital certificates, tokens and no authentication

The ASA 5505 supports preshared keys, token-based (e.g., SDI one-time passwords), and “no user authentication” for user authentication. **NOTE:** The Cisco Easy VPN server can use the digital certificate as part of user authorization. See [Chapter 27, “Configuring IPsec and ISAKMP”](#) for instructions.



CHAPTER 35

Configuring the PPPoE Client

This section describes how to configure the PPPoE client provided with the security appliance. It includes the following topics:

- [PPPoE Client Overview, page 35-1](#)
- [Configuring the PPPoE Client Username and Password, page 35-2](#)
- [Enabling PPPoE, page 35-3](#)
- [Using PPPoE with a Fixed IP Address, page 35-3](#)
- [Monitoring and Debugging the PPPoE Client, page 35-4](#)
- [Using Related Commands, page 35-5](#)

PPPoE Client Overview

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.

**Note**

PPPoE is not supported when failover is configured on the security appliance, or in multiple context or transparent mode. PPPoE is only supported in single, routed mode, without failover.

Configuring the PPPoE Client Username and Password

To configure the username and password used to authenticate the security appliance to the access concentrator, use the **vpdn** command. To use the **vpdn** command, you first define a VPDN group and then create individual users within the group.

To configure a PPPoE username and password, perform the following steps:

- Step 1** Define the VPDN group to be used for PPPoE using the following command:

```
hostname(config)# vpdn group group_name request dialout pppoe
```

In this command, replace *group_name* with a descriptive name for the group, such as “pppoe-sbc.”

- Step 2** If your ISP requires authentication, select an authentication protocol by entering the following command:

```
hostname(config)# vpdn group group_name ppp authentication {chap | mschap | pap}
```

Replace *group_name* with the same group name you defined in the previous step. Enter the appropriate keyword for the type of authentication used by your ISP:

- CHAP—Challenge Handshake Authentication Protocol
- MS-CHAP—Microsoft Challenge Handshake Authentication Protocol Version 1
- PAP—Password Authentication Protocol

**Note**

When using CHAP or MS-CHAP, the username may be referred to as the remote system name, while the password may be referred to as the CHAP secret.

- Step 3** Associate the username assigned by your ISP to the VPDN group by entering the following command:

```
hostname(config)# vpdn group group_name localname username
```

Replace *group_name* with the VPDN group name and *username* with the username assigned by your ISP.

- Step 4** Create a username and password pair for the PPPoE connection by entering the following command:

```
hostname(config)# vpdn username username password password [store-local]
```

Replace *username* with the username and *password* with the password assigned by your ISP.

**Note**

The **store-local** option stores the username and password in a special location of NVRAM on the security appliance. If an Auto Update Server sends a **clear config** command to the security appliance and the connection is then interrupted, the security appliance can read the username and password from NVRAM and re-authenticate to the Access Concentrator.

Enabling PPPoE



Note

You must complete the configuration using the **vpdn** command, described in “[Configuring the PPPoE Client Username and Password](#),” before enabling PPPoE.

The PPPoE client functionality is turned off by default. To enable PPPoE, perform the following steps:

- Step 1** Enable the PPPoE client by entering the following command from interface configuration mode:

```
hostname(config-if)# ip address pppoe [setroute]
```

The **setroute** option sets the default routes when the PPPoE client has not yet established a connection. When using the **setroute** option, you cannot have a statically defined route in the configuration.

PPPoE is not supported in conjunction with DHCP because with PPPoE the IP address is assigned by PPP. The **setroute** option causes a default route to be created if no default route exists. The default router is the address of the access concentrator. The maximum transmission unit (MTU) size is automatically set to 1492 bytes, which is the correct value to allow PPPoE transmission within an Ethernet frame.

Reenter this command to reset the DHCP lease and request a new lease.



Note

If PPPoE is enabled on two interfaces (such as a primary and backup interface), and you do not configure dual ISP support (see the “[Configuring Static Route Tracking](#)” section on page 9-5), then the security appliance can only send traffic through the first interface to acquire an IP address.

For example:

```
hostname(config)# interface gigabitethernet 0/0
hostname(config-if)# ip address pppoe
```

- Step 2** Specify a VPDN group for the PPPoE client to use with the following command from interface configuration mode (optional):

```
hostname(config-if)# pppoe client vpdn group grpname
```

grpname is the name of a VPDN group.



Note

If you have multiple VPDN groups configured, and you do not specify a group with the **pppoe client vpdn group** command, the security appliance may randomly choose a VPDN group. To avoid this, specify a VPDN group.

Using PPPoE with a Fixed IP Address

You can also enable PPPoE by manually entering the IP address, using the **ip address** command from interface configuration mode in the following format:

```
hostname(config-if)# ip address ipaddress mask pppoe
```

This command causes the security appliance to use the specified address instead of negotiating with the PPPoE server to assign an address dynamically. Replace *ipaddress* and *mask* with the IP address and subnet mask assigned to your security appliance.

For example:

```
hostname(config-if)# ip address outside 201.n.n.n 255.255.255.0 pppoe
```



Note

The **setroute** option is an option of the **ip address** command that you can use to allow the access concentrator to set the default routes when the PPPoE client has not yet established a connection. When using the **setroute** option, you cannot have a statically defined route in the configuration.

Monitoring and Debugging the PPPoE Client

Use the following command to display the current PPPoE client configuration information:

```
hostname# show ip address outside pppoe
```

Use the following command to enable or disable debugging for the PPPoE client:

```
hostname# [no] debug pppoe {event | error | packet}
```

The following summarizes the function of each keyword:

- **event**—Displays protocol event information
- **error**—Displays error messages
- **packet**—Displays packet information

Use the following command to view the status of PPPoE sessions:

```
hostname# show vpdn session [l2tp | pppoe] [id sess_id | packets | state | window]
```

The following example shows a sample of information provided by this command:

```
hostname# show vpdn

Tunnel id 0, 1 active sessions
    time since change 65862 secs
    Remote Internet Address 10.0.0.1
    Local Internet Address 199.99.99.3
    6 packets sent, 6 received, 84 bytes sent, 0 received
Remote Internet Address is 10.0.0.1
    Session state is SESSION_UP
    Time since event change 65865 secs, interface outside
    PPP interface id is 1
    6 packets sent, 6 received, 84 bytes sent, 0 received
hostname#
hostname# show vpdn session
PPPoE Session Information (Total tunnels=1 sessions=1)
Remote Internet Address is 10.0.0.1
    Session state is SESSION_UP
    Time since event change 65887 secs, interface outside
    PPP interface id is 1
    6 packets sent, 6 received, 84 bytes sent, 0 received
hostname#
hostname# show vpdn tunnel
PPPoE Tunnel Information (Total tunnels=1 sessions=1)
Tunnel id 0, 1 active sessions
    time since change 65901 secs
    Remote Internet Address 10.0.0.1
```

```
Local Internet Address 199.99.99.3
 6 packets sent, 6 received, 84 bytes sent, 0 received
hostname#
```

Clearing the Configuration

To remove all **vpdn group** commands from the configuration, use the **clear configure vpdn group** command in global configuration mode:

```
hostname(config)# clear configure vpdn group
```

To remove all **vpdn username** commands, use the **clear configure vpdn username** command:

```
hostname(config)# clear configure vpdn username
```

Entering either of these commands has no affect upon active PPPoE connections.

Using Related Commands

Use the following command to cause the DHCP server to use the WINS and DNS addresses provided by the access concentrator as part of the PPP/IPCP negotiations:

```
hostname(config)# dhcpcd auto_config [client_ifx_name]
```

This command is only required if the service provider provides this information as described in RFC 1877. The *client_ifx_name* parameter identifies the interface supported by the DHCP **auto_config** option. At this time, this keyword is not required because the PPPoE client is only supported on a single outside interface.



CHAPTER 36

Configuring LAN-to-LAN IPsec VPNs

A LAN-to-LAN VPN connects networks in different geographic locations.



Note

The ASA supports LAN-to-LAN IPsec connections with Cisco peers, and with third-party peers that comply with all relevant standards.

This chapter describes how to build a LAN-to-LAN VPN connection. It includes the following sections:

- [Summary of the Configuration, page 36-1](#)
- [Configuring Interfaces, page 36-2](#)
- [Configuring ISAKMP Policy and Enabling ISAKMP on the Outside Interface, page 36-2](#)
- [Creating a Transform Set, page 36-4](#)
- [Configuring an ACL, page 36-4](#)
- [Defining a Tunnel Group, page 36-5](#)
- [Creating a Crypto Map and Applying It To an Interface, page 36-6](#)

Summary of the Configuration

This section provides a summary of the example LAN-to-LAN configuration this chapter creates. Later sections provide step-by-step instructions.

```
hostname(config)# interface ethernet0
hostname(config-if)# ip address 10.10.4.100 255.255.0.0
hostname(config-if)# no shutdown
hostname(config)# isakmp policy 1 authentication pre-share
hostname(config)# isakmp policy 1 encryption 3des
hostname(config)# isakmp policy 1 hash sha
hostname(config)# isakmp policy 1 group 2
hostname(config)# isakmp policy 1 lifetime 43200
hostname(config)# isakmp enable outside
hostname(config)# crypto ipsec transform set FirstSet esp-3des esp-md5-hmac
hostname(config)# access-list 121_list extended permit ip 192.168.0.0 255.255.0.0
150.150.0.0 255.255.0.0
hostname(config)# tunnel-group 10.10.4.108 type ipsec-121
hostname(config)# tunnel-group 10.10.4.108 ipsec-attributes
hostname(config-ipsec)# pre-shared-key 44kkao159636jnfX
hostname(config)# crypto map abcmap 1 match address 121_list
hostname(config)# crypto map abcmap 1 set peer 10.10.4.108
hostname(config)# crypto map abcmap 1 set transform-set FirstSet
```

```
hostname(config)# crypto map abcmap interface outside
hostname(config)# write memory
```

Configuring Interfaces

A security appliance has at least two interfaces, referred to here as outside and inside. Typically, the outside interface is connected to the public Internet, while the inside interface is connected to a private network and is protected from public access.

To begin, configure and enable two interfaces on the security appliance. Then, assign a name, IP address and subnet mask. Optionally, configure its security level, speed, and duplex operation on the security appliance.

To configure interfaces, perform the following steps, using the command syntax in the examples:

-
- Step 1** To enter Interface configuration mode, in global configuration mode enter the **interface** command with the default name of the interface to configure. In the following example the interface is ethernet0.

```
hostname(config)# interface ethernet0
hostname(config-if)#
```

- Step 2** To set the IP address and subnet mask for the interface, enter the **ip address** command. In the following example the IP address is 10.10.4.100 and the subnet mask is 255.255.0.0.

```
hostname(config-if)# ip address 10.10.4.100 255.255.0.0
hostname(config-if)#
```

- Step 3** To name the interface, enter the **nameif** command, maximum of 48 characters. You cannot change this name after you set it. In the following example the name of the ethernet0 interface is outside.

```
hostname(config-if)# nameif outside
hostname(config-if)##
```

- Step 4** To enable the interface, enter the **no** version of the **shutdown** command. By default, interfaces are disabled.

```
hostname(config-if)# no shutdown
hostname(config-if)#
```

- Step 5** To save your changes, enter the **write memory** command.

```
hostname(config-if)# write memory
hostname(config-if)#
```

- Step 6** To configure a second interface, use the same procedure.
-

Configuring ISAKMP Policy and Enabling ISAKMP on the Outside Interface

The Internet Security Association and Key Management Protocol, also called IKE, is the negotiation protocol that lets two hosts agree on how to build an IPsec security association. Each ISAKMP negotiation is divided into two sections called Phase1 and Phase 2.

Phase 1 creates the first tunnel, which protects later ISAKMP negotiation messages. Phase 2 creates the tunnel that protects data travelling across the secure connection.

To set the terms of the ISAKMP negotiations, you create an ISAKMP policy, which includes the following:

- An authentication method, to ensure the identity of the peers.
- An encryption method, to protect the data and ensure privacy.
- A Hashed Message Authentication Codes method to ensure the identity of the sender and to ensure that the message has not been modified in transit.
- A Diffie-Hellman group to establish the strength of the encryption-key-determination algorithm. The security appliance uses this algorithm to derive the encryption and hash keys.
- A time limit for how long the security appliance uses an encryption key before replacing it.

See [on page 27-3](#) in the “Configuring IPsec and ISAKMP” chapter of this guide for detailed information about the IKE policy keywords and their values.

To configure ISAKMP policies, in global configuration mode use the **isakmp policy** command with its various arguments. The syntax for ISAKMP policy commands is as follows:

isakmp policy *priority* **attribute_name** [**attribute_value** | *integer*].

Perform the following steps and use the command syntax in the following examples as a guide.

-
- Step 1** Set the authentication method. The following example configures a preshared key. The priority is 1 in this and all following steps.

```
hostname(config)# isakmp policy 1 authentication pre-share
hostname(config)#
```

- Step 2** Set the encryption method. The following example configures 3DES.

```
hostname(config)# isakmp policy 1 encryption 3des
hostname(config)#
```

- Step 3** Set the HMAC method. The following example configures SHA-1.

```
hostname(config)# isakmp policy 1 hash sha
hostname(config)#
```

- Step 4** Set the Diffie-Hellman group. The following example configures Group 2.

```
hostname(config)# isakmp policy 1 group 2
hostname(config)#
```

- Step 5** Set the encryption key lifetime. The following example configures 43,200 seconds (12 hours).

```
hostname(config)# isakmp policy 1 lifetime 43200
hostname(config)#
```

- Step 6** Enable ISAKMP on the interface named outside.

```
hostname(config)# isakmp enable outside
hostname(config)#
```

- Step 7** To save your changes, enter the **write memory** command.

```
hostname(config)# write memory
hostname(config)#
```

Creating a Transform Set

A transform set combines an encryption method and an authentication method. During the IPsec security association negotiation with ISAKMP, the peers agree to use a particular transform set to protect a particular data flow. The transform set must be the same for both peers.

A transform set protects the data flows for the access list specified in the associated crypto map entry. You can create transform sets in the security appliance configuration, and then specify a maximum of 11 of them in a crypto map or dynamic crypto map entry.

Table 36-1 lists valid encryption and authentication methods.

Table 36-1 Valid Encryption and Authentication Methods

Valid Encryption Methods	Valid Authentication Methods
esp-des	esp-md5-hmac
esp-3des (default)	esp-sha-hmac (default)
esp-aes (128-bit encryption)	
esp-aes-192	
esp-aes-256	
esp-null	

Tunnel Mode is the usual way to implement IPsec between two security appliances that are connected over an untrusted network, such as the public Internet. Tunnel mode is the default and requires no configuration.

To configure a transform set, perform the following steps:

- Step 1** In global configuration mode enter the **crypto ipsec transform-set** command. The following example configures a transform set with the name FirstSet, esp-3des encryption, and esp-md5-hmac authentication. The syntax is as follows:

crypto ipsec transform-set *transform-set-name encryption-method authentication-method*

```
hostname(config)# crypto ipsec transform-set FirstSet esp-3des esp-md5-hmac
hostname(config)#
```

- Step 2** Save your changes.

```
hostname(config)# write memory
hostname(config)#
```

Configuring an ACL

The security appliance uses access control lists to control network access. By default, the security appliance denies all traffic. You need to configure an ACL that permits traffic.

The ACLs that you configure for this LAN-to-LAN VPN control connections are based on the source and destination IP addresses. Configure ACLs that mirror each other on both sides of the connection.

To configure an ACL, perform the following steps:

- Step 1** Enter the **access-list extended** command. The following example configures an ACL named `l2l_list` that lets traffic from IP addresses in the 192.168.0.0 network travel to the 150.150.0.0 network. The syntax is **access-list listname extended permit ip source-ipaddress source-netmask destination-ipaddress destination-netmask**.

```
hostname(config)# access-list l2l_list extended permit ip 192.168.0.0 255.255.0.0
150.150.0.0 255.255.0.0
hostname(config)#
```

- Step 2** Configure an ACL for the security appliance on the other side of the connection that mirrors the ACL above. In the following example the prompt for the peer is `hostname2`.

```
hostname2(config)# access-list l2l_list extended permit ip 150.150.0.0 255.255.0.0
192.168.0.0 255.255.0.0
hostname2(config)#
```

Defining a Tunnel Group

A tunnel group is a set of records that contain tunnel connection policies. You configure a tunnel group to identify AAA servers, specify connection parameters, and define a default group policy. The security appliance stores tunnel groups internally.

There are two default tunnel groups in the security appliance system: `DefaultRAGroup`, which is the default IPsec remote-access tunnel group, and `DefaultL2Lgroup`, which is the default IPsec LAN-to-LAN tunnel group. You can modify them but not delete them. You can also create one or more new tunnel groups to suit your environment. The security appliance uses these groups to configure default tunnel parameters for remote access and LAN-to-LAN tunnel groups when there is no specific tunnel group identified during tunnel negotiation.

To establish a basic LAN-to-LAN connection, you must set two attributes for a tunnel group:

- Set the connection type to IPsec LAN-to-LAN.
- Configure an authentication method, in the following example, preshared key.

- Step 1** To set the connection type to IPsec LAN-to-LAN, enter the **tunnel-group** command. The syntax is **tunnel-group name type type**, where *name* is the name you assign to the tunnel group, and *type* is the type of tunnel. The tunnel types as you enter them in the CLI are:

- **ipsec-ra** (IPsec remote access)
- **ipsec-l2l** (IPsec LAN to LAN)

In the following example the name of the tunnel group is the IP address of the LAN-to-LAN peer, 10.10.4.108.

```
hostname(config)# tunnel-group 10.10.4.108 type ipsec-l2l
hostname(config)#
```

- Step 2** To set the authentication method to preshared key, enter the `ipsec-attributes` mode and then enter the **pre-shared-key** command to create the preshared key. You need to use the same preshared key on both security appliances for this LAN-to-LAN connection.

The key is an alphanumeric string of 1-128 characters. In the following example the preshared key is 44kkaol59636jnfx.

```
hostname(config)# tunnel-group 10.10.4.108 ipsec-attributes
```

```
hostname(config-ipsec)# pre-shared-key 44kkao159636jnfx
```

Step 3 Save your changes.

```
hostname(config)# write memory  
hostname(config)#
```

Creating a Crypto Map and Applying It To an Interface

Crypto map entries pull together the various elements of IPsec security associations, including the following:

- Which traffic IPsec should protect, which you define in an access list.
- Where to send IPsec-protected traffic, by identifying the peer.
- What IPsec security applies to this traffic, which a transform set specifies.
- The local address for IPsec traffic, which you identify by applying the crypto map to an interface.

For IPsec to succeed, both peers must have crypto map entries with compatible configurations. For two crypto map entries to be compatible, they must, at a minimum, meet the following criteria:

- The crypto map entries must contain compatible crypto access lists (for example, mirror image access lists). If the responding peer uses dynamic crypto maps, the entries in the security appliance crypto access list must be “permitted” by the peer’s crypto access list.
- The crypto map entries each must identify the other peer (unless the responding peer is using a dynamic crypto map).
- The crypto map entries must have at least one transform set in common.

If you create more than one crypto map entry for a given interface, use the sequence number (seq-num) of each entry to rank it: the lower the seq-num, the higher the priority. At the interface that has the crypto map set, the security appliance evaluates traffic against the entries of higher priority maps first.

Create multiple crypto map entries for a given interface if either of the following conditions exist:

- Different peers handle different data flows.
- You want to apply different IPsec security to different types of traffic (to the same or separate peers), for example, if you want traffic between one set of subnets to be authenticated, and traffic between another set of subnets to be both authenticated and encrypted. In this case, define the different types of traffic in two separate access lists, and create a separate crypto map entry for each crypto access list.

To create a crypto map and apply it to the outside interface in global configuration mode, enter several of the **crypto map** commands. These commands use a variety of arguments, but the syntax for all of them begin with **crypto map map-name-seq-num**. In the following example the map-name is **abcmmap**, the sequence number is 1.

Enter these commands in global configuration mode:

Step 1 To assign an access list to a crypto map entry, enter the **crypto map match address** command.

The syntax is **crypto map map-name seq-num match address aclname**. In the following example the map name is **abcmmap**, the sequence number is 1, and the access list name is **121_list**.

```
hostname(config)# crypto map abcmmap 1 match address 121_list  
hostname(config)#
```

Step 2 To identify the peer (s) for the IPsec connection, enter the **crypto map set peer** command.

The syntax is **crypto map map-name seq-num set peer** {*ip_address1* | *hostname1*}[... *ip_address10* | *hostname10*]. In the following example the peer name is 10.10.4.108.

```
hostname(config)# crypto map abcmmap 1 set peer 10.10.4.108  
hostname(config)#
```

Step 3 To specify a transform set for a crypto map entry, enter the **crypto map set transform-set** command.

The syntax is **crypto map map-name seq-num set transform-set transform-set-name**. In the following example the transform set name is FirstSet.

```
hostname(config)# crypto map abcmmap 1 set transform-set FirstSet  
hostname(config)#
```

Applying Crypto Maps to Interfaces

You must apply a crypto map set to each interface through which IPsec traffic travels. The security appliance supports IPsec on all interfaces. Applying the crypto map set to an interface instructs the security appliance to evaluate all interface traffic against the crypto map set and to use the specified policy during connection or security association negotiations.

Binding a crypto map to an interface also initializes the runtime data structures, such as the security association database and the security policy database. When you later modify a crypto map in any way, the security appliance automatically applies the changes to the running configuration. It drops any existing connections and reestablishes them after applying the new crypto map.

Step 1 To apply the configured crypto map to the outside interface, enter the **crypto map interface** command. The syntax is **crypto map map-name interface interface-name**.

```
hostname(config)# crypto map abcmmap interface outside  
hostname(config)#
```

Step 2 Save your changes.

```
hostname(config)# write memory  
hostname(config)#
```



CHAPTER 37

Configuring WebVPN

This chapter includes the following sections:

- [Getting Started with WebVPN, page 37-1](#)
- [Creating and Applying WebVPN Policies, page 37-15](#)
- [Configuring WebVPN Tunnel Group Attributes, page 37-17](#)
- [Configuring WebVPN Group Policy and User Attributes, page 37-17](#)
- [Configuring Application Access, page 37-18](#)
- [Configuring File Access, page 37-22](#)
- [Configuring Access to Citrix MetaFrame Services, page 37-24](#)
- [Using WebVPN with PDAs, page 37-25](#)
- [Using E-Mail over WebVPN, page 37-26](#)
- [Optimizing WebVPN Performance, page 37-28](#)
- [WebVPN End User Setup, page 37-32](#)
- [Capturing WebVPN Data, page 37-50](#)

Getting Started with WebVPN

WebVPN lets users establish a secure, remote-access VPN tunnel to a security appliance using a web browser. Users do not need a software or hardware client.

WebVPN provides secure and easy access to a broad range of web resources and web-enabled applications from almost any computer on the Internet. They include:

- Internal websites
- Web-enabled applications
- NT/Active Directory file shares
- E-mail proxies, including POP3S, IMAP4S, and SMTPS
- MS Outlook Web Access
- MAPI
- Application Access (that is, port forwarding for access to other TCP-based applications)

WebVPN uses Secure Sockets Layer Protocol and its successor, Transport Layer Security to provide the secure connection between remote users and specific, supported internal resources that you configure at a central site. The security appliance recognizes connections that need to be proxied, and the HTTP server interacts with the authentication subsystem to authenticate users.

The network administrator provides access to WebVPN resources to users on a group basis. Users have no direct access to resources on the internal network.

The following sections address getting started with the configuration of WebVPN access:

- [Observing WebVPN Security Precautions](#)
- [Understanding Features Not Supported for WebVPN](#)
- [Using SSL to Access the Central Site](#)
- [Authenticating with Digital Certificates](#)
- [Enabling Cookies on Browsers for WebVPN](#)
- [Managing Passwords](#)
- [Using Single Sign-on with WebVPN](#)
- [Authenticating with Digital Certificates](#)

Observing WebVPN Security Precautions

WebVPN connections on the security appliance are very different from remote access IPSec connections, particularly with respect to how they interact with SSL-enabled servers, and precautions to reduce security risks.

In a WebVPN connection, the security appliance acts as a proxy between the end user web browser and target web servers. When a WebVPN user connects to an SSL-enabled web server, the security appliance establishes a secure connection and validates the server SSL certificate. The end user browser never receives the presented certificate, so therefore cannot examine and validate the certificate.

The current implementation of WebVPN on the security appliance does not permit communication with sites that present expired certificates. Nor does the security appliance perform trusted CA certificate validation. Therefore, WebVPN users cannot analyze the certificate an SSL-enabled web-server presents before communicating with it.

To minimize the risks involved with SSL certificates:

1. Configure a group policy that consists of all users who need WebVPN access and enable the WebVPN feature only for that group policy.
2. Limit Internet access for WebVPN users. One way to do this is to disable URL entry. Then configure links to specific targets within the private network that you want WebVPN users to be able to access.
3. Educate users. If an SSL-enabled site is not inside the private network, users should not visit this site over a WebVPN connection. They should open a separate browser window to visit such sites, and use that browser to view the presented certificate.

Understanding Features Not Supported for WebVPN

The security appliance does not support the following features for WebVPN connections:

- Inspection features under the Modular Policy Framework, inspecting configuration control.
- Functionality the filter configuration commands provide, including the **vpn-filter** command.

- NAT, reducing the need for globally unique IP addresses.
- PAT, permitting multiple outbound sessions appear to originate from a single IP address.
- QoS, rate limiting using the **police** command and **priority-queue** command.
- Connection limits, checking either via the static or the Modular Policy Framework **set connection** command.
- The **established** command, allowing 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.

Using SSL to Access the Central Site

WebVPN uses SSL and its successor, TLS1 to provide a secure connection between remote users and specific, supported internal resources at a central site. This section includes the following topics:

- [Using HTTPS for WebVPN Sessions](#)
- [Configuring WebVPN and ASDM on the Same Interface](#)
- [Setting WebVPN HTTP/HTTPS Proxy](#)
- [Configuring SSL/TLS Encryption Protocols](#)

Using HTTPS for WebVPN Sessions

Establishing WebVPN sessions requires the following:

- Using HTTPS to access the security appliance or load balancing cluster. In a web browser, users enter the security appliance IP address in the format *https:// address* where *address* is the IP address or DNS hostname of the security appliance interface.
- Enabling WebVPN sessions on the security appliance interface that users connect to.

To permit WebVPN sessions on an interface, perform the following steps:

-
- Step 1** In global configuration mode, enter the **webvpn** command to enter webvpn mode.
- Step 2** Enter the **enable** command with the name of the interface that you want to use for WebVPN sessions.

For example, to enable WebVPN sessions on the interface called outside, enter the following:

```
hostname(config)# webvpn
hostname(config-webvpn)# enable outside
```

Configuring WebVPN and ASDM on the Same Interface

The security appliance can support both WebVPN connections and HTTPS connections for ASDM administrative sessions simultaneously on the same interface. Both HTTPS and WebVPN use port 443 by default. Therefore, to enable both HTTPS and WebVPN on the same interface, you must specify a different port number for either HTTPS or WebVPN. An alternative is to configure WebVPN and HTTPS on different interfaces.

To specify a port for HTTPS, use the *port* argument of the **http server enable** command. The following example specifies that HTTPS ASDM sessions use port 444 on the outside interface. WebVPN is also enabled on the outside interface and uses the default port (443). With this configuration, remote users initiate ASDM sessions by entering `https://<outside_ip>:444` in the browser.

```
hostname(config)# http server enable 444
hostname(config)# http 192.168.3.0 255.255.255.0 outside
hostname(config)# webvpn
hostname(config-webvpn)# enable outside
```

To specify a port for WebVPN, use the **port** command from webvpn configuration mode. The next example enables WebVPN on port 444 of the outside interface. HTTPS for ASDM is also configured on the outside interface and uses the default port (443). With this configuration, remote users initiating WebVPN sessions enter `https://<outside_ip>:444` in the browser.

```
hostname(config)# http server enable
hostname(config)# http 192.168.3.0 255.255.255.0 outside
hostname(config)# webvpn
hostname(config-webvpn)# port 444
hostname(config-webvpn)# enable outside
```

Setting WebVPN HTTP/HTTPS Proxy

The security appliance can terminate HTTPS connections and forward HTTP/HTTPS requests to HTTP and HTTPS proxy servers. These servers act as intermediaries between users and the Internet. Requiring all Internet access via a server that the organization controls provides another opportunity for filtering to assure secure Internet access and administrative control.

To set values for HTTP and HTTPS proxy, use the **http-proxy** and **https-proxy** commands in webvpn mode. These commands let you identify HTTP and HTTPS proxy servers and ports.

Configuring SSL/TLS Encryption Protocols

When you set SSL/TLS encryption protocols, be aware of the following:

- Make sure that the security appliance and the browser you use allow the same SSL/TLS encryption protocols.
- If you configure e-mail proxy, do not set the security appliance SSL version to TLSv1 Only. MS Outlook and MS Outlook Express do not support TLS.
- TCP Port Forwarding requires Sun Microsystems Java Runtime Environment (JRE) version 1.4.x and 1.5.x. Port forwarding does not work when a WebVPN user connects with some SSL versions, as follows:

Negotiate SSLv3	Java downloads
Negotiate SSLv3/TLSv1	Java downloads
Negotiate TLSv1	Java does NOT download
TLSv1Only	Java does NOT download
SSLv3Only	Java does NOT download

Authenticating with Digital Certificates

SSL uses digital certificates for authentication. The security appliance creates a self-signed SSL server certificate when it boots; or you can install in the security appliance an SSL certificate that has been issued in a PKI context. For HTTPS, this certificate must then be installed on the client. You need to install the certificate from a given security appliance only once.

Restrictions for authenticating users with digital certificates include the following:

- Application Access does not work for WebVPN users who authenticate using digital certificates. JRE does not have the ability to access the web browser keystore. Therefore JAVA cannot use a certificate that the browser uses to authenticate a user, so it cannot start.
- E-mail proxy supports certificate authentication with Netscape 7.x e-mail clients only. Other e-mail clients such as MS Outlook, MS Outlook Express, and Eudora lack the ability to access the certificate store.

For more information on authentication and authorization using digital certificates, see “[Using Certificates and User Login Credentials](#)” in the “[Configuring AAA Servers and the Local Database](#)” chapter.

Enabling Cookies on Browsers for WebVPN

Browser cookies are required for the proper operation of WebVPN. When cookies are disabled on the web browser, the links from the web portal home page open a new window prompting the user to log in once more.

Managing Passwords

You can configure the security appliance to warn end users when their passwords are about to expire. To do this, you specify the **password-management** command in tunnel-group general-attributes mode.

When you configure this command, the security appliance notifies the remote user at login that the user’s current password is about to expire or has expired. The security appliance then offers the user the opportunity to change the password. If the current password has not yet expired, the user can still log in using that password. This command is valid for AAA servers that support such notification; that is, RADIUS, RADIUS with an NT server, and LDAP servers. The security appliance ignores this command if RADIUS or LDAP authentication has not been configured.

Note that this does not change the number of days before the password expires, but rather specifies the number of days ahead of expiration that the security appliance starts warning the user that the password is about to expire. The default value is 14 days.

For LDAP server authentication only, you can use the **password-expire-in-days** keyword to specify a specific number of days. If you specify the **password-expire-in-days** keyword, you must also specify the number of days.

Specifying this command with the number of days set to 0 disables this command. The security appliance then does not notify the user of the pending expiration, but the user can change the password after it expires.

The following example sets the days before password expiration to begin warning the user of the pending expiration to 90 for the tunnel group “testgroup”:

```
hostname(config)# tunnel-group testgroup type webvpn
hostname(config)# tunnel-group testgroup general-attributes
```

```
hostname(config-general)# password-management password-expire-in-days 90
```

Using Single Sign-on with WebVPN

Single sign-on support lets WebVPN users enter a username and password only once to access multiple protected services and web servers. In general, the SSO mechanism either starts as part of the AAA process or just after successful user authentication to a AAA server. The WebVPN server running on the security appliance acts as a proxy for the user to the authenticating server. When a user logs in, the WebVPN server sends an SSO authentication request, including username and password, to the authenticating server using HTTPS. If the server approves the authentication request, it returns an SSO authentication cookie to the WebVPN server. The security appliance keeps this cookie on behalf of the user and uses it to authenticate the user to secure websites within the domain protected by the SSO server.

This section describes the three SSO authentication methods supported by WebVPN: HTTP Basic and NTLMv1 (NT LAN Manager) authentication, the Computer Associates eTrust SiteMinder SSO server (formerly Netegrity SiteMinder), and the HTTP Form protocol.

This section includes:

- [Configuring SSO with HTTP Basic or NTLM Authentication](#)
- [Configuring SSO Authentication Using SiteMinder](#)
- [Configuring SSO with the HTTP Form Protocol](#)

Configuring SSO with HTTP Basic or NTLM Authentication

This section describes single sign-on with HTTP Basic or NTLM authentication. You can configure the security appliance to implement SSO using either or both of these methods. The **auto-signon** command configures the security appliance to automatically pass WebVPN user login credentials (username and password) on to internal servers. You can enter multiple **auto-signon** commands. The security appliance processes them according to the input order (early commands take precedence). You specify the servers to receive the login credentials using either IP address and IP mask, or URI mask.

Use the **auto-signon** command in any of three modes: webvpn configuration, webvpn group-policy mode, or webvpn username mode. Username supersedes group, and group supersedes global. The mode you choose depends upon scope of authentication you want:

Mode	Scope
Webvpn configuration	All WebVPN users globally
Webvpn group configuration	A subset of WebVPN users defined by a group policy
Webvpn username configuration	An individual WebVPN user

The following example commands present various possible combinations of modes and arguments.

All Users, IP Address Range, NTLM

To configure **auto-signon** for all WebVPN users to servers with IP addresses ranging from 10.1.1.0 to 10.1.1.255 using NTLM authentication, for example, enter the following commands:

```
hostname(config)# webvpn
hostname(config-webvpn)# auto-signon allow ip 10.1.1.1 255.255.255.0 auth-type ntlm
```

All Users, URI Range, HTTP Basic

To configure **auto-signon** for all WebVPN users, using basic HTTP authentication, to servers defined by the URI mask `https://*.example.com/*`, for example, enter the following commands:

```
hostname(config)# webvpn
hostname(config-webvpn)# auto-signon allow uri https://*.example.com/* auth-type basic
```

Group, URI Range, HTTP Basic and NTLM

To configure **auto-signon** for WebVPN users ExamplePolicy group policy, using either basic or NTLM authentication, to servers defined by the URI mask `https://*.example.com/*`, for example, enter the following commands:

```
hostname(config)# group-policy ExamplePolicy attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# auto-signon allow uri https://*.example.com/* auth-type all
```

Specific User, IP Address Range, HTTP Basic

To configure **auto-signon** for a user named Anyuser to servers with IP addresses ranging from 10.1.1.0 to 10.1.1.255 using HTTP Basic authentication, for example, enter the following commands:

```
hostname(config)# username Anyuser attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# auto-signon allow ip 10.1.1.1 255.255.255.0 auth-type basic
```

Configuring SSO Authentication Using SiteMinder

This section describes configuring the security appliance to support SSO with SiteMinder. You would typically choose to implement SSO with SiteMinder if your website security infrastructure already incorporates SiteMinder. With this method, SSO authentication is separate from AAA and happens once the AAA process completes. If you want to configure SSO for a WebVPN user or group, you must first configure a AAA server, such as a RADIUS or LDAP server. You can then setup SSO support for WebVPN. This section includes:

- [Task Overview: Configuring SSO with Siteminder](#)
- [Detailed Tasks: Configuring SSO with Siteminder](#)
- [Adding the Cisco Authentication Scheme to SiteMinder](#)

Task Overview: Configuring SSO with Siteminder

This section presents an overview of the tasks necessary to configure SSO with SiteMinder SSO. These tasks are:

- Specifying the SSO server.
- Specifying the URL of the SSO server to which the security appliance makes SSO authentication requests.
- Specifying a secret key to secure the communication between the security appliance and the SSO server. This key is similar to a password: you create it, save it, and enter it on both the security appliance and the SiteMinder Policy Server using the Cisco Java plug-in authentication scheme.

In addition to these required tasks, you can optionally do the following configuration tasks:

- Configuring the authentication request timeout.

- Configuring the number of authentication request retries.

After you have completed the configuration tasks, you assign an SSO server to a user or group policy.

Detailed Tasks: Configuring SSO with Siteminder

This section presents specific steps for configuring the security appliance to support SSO authentication with CA SiteMinder. To configure SSO with SiteMinder, perform the following steps:

- Step 1** In webvpn configuration mode, enter the **sso-server** command with the **type** option to create an SSO server. For example, to create an SSO server named Example of type siteminder, enter the following:

```
hostname(config)# webvpn
hostname(config-webvpn)# sso-server Example type siteminder
hostname(config-webvpn-sso-siteminder)#
```



Note

The security appliance currently supports only the SSO server type siteminder.

- Step 2** Enter the **web-agent-url** command in webvpn-sso-siteminder configuration mode to specify the authentication URL of the SSO server. For example, to send authentication requests to the URL <http://www.Example.com/webvpn>, enter the following:

```
hostname(config-webvpn-sso-siteminder)# web-agent-url http://www.Example.com/webvpn
hostname(config-webvpn-sso-siteminder)#
```

- Step 3** Specify a secret key to secure the authentication communications between the security appliance and SiteMinder using the **policy-server-secret** command in webvpn-sso-siteminder configuration mode. You can create a key of any length using any regular or shifted alphanumeric character, but you must enter the same key on both the security appliance and the SSO server.

For example, to create the secret key AtaL8rD8!, enter the following:

```
hostname(config-webvpn-sso-siteminder)# policy-server-secret AtaL8rD8!
hostname(config-webvpn-sso-siteminder)#
```

- Step 4** Optionally, you can configure the number of seconds before a failed SSO authentication attempt times out using the **request-timeout** command in webvpn-sso-siteminder configuration mode. The default number of seconds is 5 seconds and the possible range is 1 to 30 seconds. To change the number of seconds before a request times out to 8, for example, enter the following:

```
hostname(config-webvpn-sso-siteminder)# request-timeout 8
hostname(config-webvpn-sso-siteminder)#
```

- Step 5** Optionally, you can configure the number of times the security appliance retries a failed SSO authentication attempt before the authentication times-out using the **max-retry-attempts** command in webvpn-sso-siteminder configuration mode. The default is 3 retry attempts and the possible range is 1 to 5 attempts. To configure the number of retries to be 4, for example, enter the following:

```
hostname(config-webvpn-sso-siteminder)# max-retry-attempts 4
hostname(config-webvpn-sso-siteminder)#
```

- Step 6** After you configure the SSO server, you must specify SSO authentication for either a group or user. To specify SSO for a group, assign an SSO server to a group policy using the **sso-server value** command in group-policy-webvpn configuration mode. To specify SSO for a user, assign an SSO server to a user policy using the same command, **sso-server value**, but in username-webvpn configuration mode. For example, to assign the SSO server named Example to the user named Anyuser, enter the following:

```
hostname(config)# username Anyuser attributes
hostname(config-username)# webvpn
```

```
hostname(config-username-webvpn)# sso-server value Example
hostname(config-group-webvpn)#
```

Step 7 Finally, you can test the SSO server configuration using the **test sso-server** command in privileged EXEC mode. For example, to test the SSO server named Example using the username Anyuser, enter the following:

```
hostname# test sso-server Example username Anyuser
INFO: Attempting authentication request to sso-server Example for user Anyuser
INFO: STATUS: Success
hostname#
```

Adding the Cisco Authentication Scheme to SiteMinder

Besides configuring the security appliance for SSO with SiteMinder, you must also configure your CA SiteMinder Policy Server with the Cisco authentication scheme, provided as a Java plug-in.



Note

- Configuring the SiteMinder Policy Server requires experience with SiteMinder.
- This section presents general tasks, not a complete procedure.
- Refer to the CA SiteMinder documentation for the complete procedure for adding a custom authentication scheme.

To configure the Cisco authentication scheme on your SiteMinder Policy Server, perform these following tasks:

Step 1 With the Siteminder Administration utility, create a custom authentication scheme being sure to use the following specific arguments:

- In the Library field, enter **smjavaapi**.
- In the Secret field, enter the same secret configured on the security appliance.
You configure this on the security appliance with either the **policy-server-secret** command at the command line interface or in the Secret Key field of the Add SSO Server dialog in ASDM.
- In the Parameter field, enter **CiscoAuthAPI**.

Step 2 Using your Cisco.com login, download the file **cisco_vpn_auth.jar** from <http://www.cisco.com/cgi-bin/tablebuild.pl/asa> and copy it to the default library directory for the SiteMinder server.

Configuring SSO with the HTTP Form Protocol

This section describes using the HTTP Form protocol for SSO. HTTP Form protocol is a common approach to SSO authentication that can also qualify as a AAA method. It provides a secure method for exchanging authentication information between WebVPN users and authenticating web servers. As a common protocol, it is highly compatible with web servers and web-based SSO products, and you can use it in conjunction with other AAA servers such as RADIUS or LDAP servers.




Note

To configure SSO with the HTTP protocol correctly, you must have a thorough working knowledge of authentication and HTTP protocol exchanges.

The security appliance again serves as a proxy for WebVPN users to an authenticating web server but, in this case, it uses HTTP Form protocol and the POST method for requests. You must configure the security appliance to send and receive form data. [Figure 37-1](#) illustrates the following SSO authentication steps:

1. A WebVPN user first enters a username and password to log into the WebVPN server on the security appliance.
2. The WebVPN server acts as a proxy for the user and forwards the form data (username and password) to an authenticating web server using a POST authentication request.
3. If the authenticating web server approves the user data, it returns an authentication cookie to the WebVPN server where it is stored on behalf of the user.
4. The WebVPN server establishes a tunnel to the user.
5. The user can now access other websites within the protected SSO environment without reentering a username and password.

Figure 37-1 SSO Authentication Using HTTP Forms



While you would expect to configure form parameters that let the security appliance include POST data such as the username and password, you initially might not be aware of additional hidden parameters that the web server requires. Some authentication applications expect hidden data which is neither visible to nor entered by the user. You can, however, discover hidden parameters the authenticating web server expects by making a direct authentication request to the web server from your browser without the security appliance in the middle acting as a proxy. Analyzing the web server response using an HTTP header analyzer reveals hidden parameters in a format similar to the following:

```
<param name>=<URL encoded value>&<param name>=<URL encoded>
```

Some hidden parameters are mandatory and some are optional. If the web server requires data for a hidden parameter, it rejects any authentication POST request that omits that data. Because a header analyzer does not tell you if a hidden parameter is mandatory or not, we recommend that you include all hidden parameters until you determine which are mandatory.

This section describes:

- [Gathering HTTP Form Data](#)
- [Task Overview: Configuring SSO with HTTP Form Protocol](#)
- [Detailed Tasks: Configuring SSO with HTTP Form Protocol](#)

Gathering HTTP Form Data

This section presents the steps for discovering and gathering necessary HTTP Form data. If you do not know what parameters the authenticating web server requires, you can gather parameter data by analyzing an authentication exchange using the following steps:



Note

These steps require a browser and an HTTP header analyzer.


- Step 1** Start your browser and HTTP header analyzer, and connect directly to the web server login page without going through the security appliance.
- Step 2** After the web server login page has loaded in your browser, examine the login sequence to determine if a cookie is being set during the exchange. If the web server has loaded a cookie with the login page, configure this login page URL as the *start-URL*.
- Step 3** Enter the username and password to log in to the web server, and press Enter. This action generates the authentication POST request that you examine using the HTTP header analyzer.

An example POST request—with host HTTP header and body—follows:

```
POST
/emco/myemco/authc/forms/MCOlogin.fcc?TYPE=33554433&REALMOID=06-000430e1-7443-125c-ac05-83
846dc90034&GUID=&SMAUTHREASON=0&METHOD=GET&SMAGENTNAME=$SM$5FZmjnk3DRNwNjk2KcqVCFbIrNT9%2b
J0H0KPshFtg6rB1UV2PxkHqLw%3d%3d&TARGET=https%3A%2F%2Fwww.example.com%2Femco%2Fmyemco%2F
HTTP/1.1
Host: www.example.com
(BODY)
SMENC=ISO-8859-1&SMLOCALE=US-EN&USERID=Anyuser&USER_PASSWORD=XXXXXX&target=https%3A%2F%2F
www.example.com%2Femco%2Fmyemco%2F&smauthreason=0
```

- Step 4** Examine the POST request and copy the protocol, host, and the complete URL to configure the action-uri parameter.
- Step 5** Examine the POST request body and copy the following:
 - a. Username parameter. In the preceding example, this parameter is USERID, not the value anyuser.
 - b. Password parameter. In the preceding example, this parameter is USER_PASSWORD.
 - c. Hidden parameter. This parameter is everything in the POST body except the username and password parameters. In the preceding example, the hidden parameter is:
 SMENC=ISO-8859-1&SMLOCALE=US-EN&target=https%3A%2F%2Fwww.example.com%2Femco%2Fmyemco%2F&smauthreason=0

Figure 37-2 highlights the action URI, hidden, username and password parameters within sample output from an HTTP analyzer. This is only an example; output varies widely across different websites.

Figure 37-2 Action-uri, hidden, username and password parameters

1	Action URI parameter
2	Hidden parameters
3	Username and password parameters

Step 6 If you successfully log in to the web server, examine the server response with the HTTP header analyzer to locate the name of the session cookie set by the server in your browser. This is the **auth-cookie-name** parameter.

In the following server response header, the name of the session cookie is SMSESSION. You just need the name, not the value.

Set-Cookie:


```
SMSESSION=yN4Yp5hHVNDgs4FT8dn7+Rwev41hse49X1Kc+1twie0gqnjbhkTkUnr8XWP3hvdH6PZPbHIHtWLDKta8
ngDB/1bYTjIxrbdX8WPWwag3CxVa3ad0xHFR8yjD55GevK3ZF4ujgU1lh06fta0dSSOSepWvnsCb7IFxCw+MGiw0o8
8uHa2t4l+SillqfJvcpuXfiIAO06D/gtDF40Ow5YKHEl2KhDEvv+yQzxwfEz2c17Ef5iMr8LgGcDK7qvMcVrgUqx68
JQOK2+RSwtHQ15bCZmsDU5vQVCvSQWC8OMHNGwps253XwRLvd/h6S/tM0k98QMv+i3N8oOdj1V7f1BqecH7+kVrU01
F6oFzr0zM1kMyLr5Hh1VDh7B0k9wp0dUFZiAza43jupD5f6CEkuLeudYWLxgNzsR8eqtPK6t1gFJyOn0s7QdNQ7q9
knsPJsekRAH9hrLBhWBLTU/3B1QS94wEGD2YTuiw36TiP14hYwOlCAYRj2/bY3+1YzVu7EmzMQ+UefYxh4cf2gYD8R
ZL2RwmP9JV5148I3XBFPNUw/3V5j7f7nRuLr/CdfK3008+Pa3V6/nNhokErSgyxjzMD88DVzM41LxxaUDhbcmkHT9I
mzBvKzJX0J+o7FoUDFOxEdIqlAN4GNqk49cpi2sXDbIarALp6B13+tbB4MLHGH+0CPscZXqoi/kon9YmGauHyRs+0m
6wthdlAmCnv1JCDfDoXtn8DpabgiW6VDTTrv13SGPyQtUv7Wdahuq5SxbUzjY2JxQnrUtwB977NCzYu2sOtN+dsEReW
```



```
J6ueyJBbMzKyzUB4L3i5uSYN50B4PCv1w5KdRka5p3N0Nfq6RM6dfipMEJw0Ny1sZ7ohz3fbvQ/YZ7lw/k7ods/8Vb
aR15ivkE8dSCzuf/AInHtCzuQ6wApzEp9CUoG8/dapWriHjNoi411JOgCst33wEhxFxcWy2UWxs4EZSjsI5GyBnefS
QTPVfma5dc/emWor9vWr0HnTQaHP5rg5dTNqunkDEdMIHfbeP3F90cZejVzihM6igiS6P/CEJAjE;
Domain=.example.com;Path=/
```

Figure 37-3 shows an example of authorization cookies in HTTP analyzer output. This is only an example; output varies widely across different websites.

Figure 37-3 Authorization cookies in sample HTTP analyzer output



1 AUTH=; path=/; domain=.example.com; expires=Wed, 01-Jan-1997 00:00:00
 1 SAUTH=; path=/; domain=.example.com; expires=Wed, 01-Jan-1997 00:00:00

1 Authorization cookies

Step 7 In some cases, the server may set the same cookie regardless of whether the authentication was successful or not, and such a cookie is unacceptable for SSO purposes. To confirm that the cookies are different, repeat [Step 1](#) through [Step 6](#) using invalid login credentials and then compare the “failure” cookie with the “success” cookie.

You now have the necessary parameter data to configure the security appliance for SSO with HTTP Form protocol.

Task Overview: Configuring SSO with HTTP Form Protocol

This section presents an overview of configuring SSO with the HTTP Form protocol. To enable SSO using HTTP Forms, perform the following tasks:

- Configure the uniform resource identifier on the authenticating web server to receive and process the form data (**action-uri**).
- Configure the username parameter (**user-parameter**).
- Configure the user password parameter (**password-parameter**).

You might also need to do the following tasks depending upon the requirements of authenticating web server:

- Configure a starting URL if the authenticating web server requires a pre-login cookie exchange (**start-url**).
- Configure any hidden authentication parameters required by the authenticating web server (**hidden-parameter**).
- Configure the name of an authentication cookie set by the authenticating web server (**auth-cookie-name**).

Detailed Tasks: Configuring SSO with HTTP Form Protocol

This section presents the detailed tasks required to configure SSO with the HTTP Form protocol. Perform the following steps to configure the security appliance to use HTTP Form protocol for SSO:

- Step 1** If the authenticating web server requires it, enter the **start-url** command in aaa-server-host configuration mode to specify the URL from which to retrieve a pre-login cookie from the authenticating web server. For example, to specify the authenticating web server URL `http://example.com/east/Area.do?Page-Grp1` in the `testgrp1` server group with an IP address of 10.0.0.2, enter the following:

```
hostname(config)# aaa-server testgrp1 host 10.0.0.2
hostname(config-aaa-server-host)# start-url http://example.com/east/Area.do?Page-Grp1
hostname(config-aaa-server-host)#
```

- Step 2** To specify a URI for an authentication program on the authenticating web server, enter the **action-uri** command in aaa-server- host configuration mode. A URI can be entered on multiple, sequential lines. The maximum number of characters per line is 255. The maximum number of characters for a complete URI is 2048. An example action URI follows:

```
http://www.example.com/auth/index.html/appdir/authc/forms/MCologin.fcc?TYPE=33554433&REALM
OID=06-000a1311-a828-1185-ab41-8333b16a0008&GUID=&SMAUTHREASON=0&METHOD=GET&SMAGENTNAME=$S
M$5FZmjnk3DRNwNjk2KcqVCFbIrNT9%2bJ0H0KpshFtg6rB1UV2PxxHqLw%3d%3d&TARGET=https%3A%2F%2Fauth
.example.com
```

To specify this action URI, enter the following commands:

```
hostname(config-aaa-server-host)# action-uri http://www.example.com/auth/index.htm
hostname(config-aaa-server-host)# action-uri 1/appdir/authc/forms/MCologin.fcc?TYP
hostname(config-aaa-server-host)# action-uri 554433&REALMOID=06-000a1311-a828-1185
hostname(config-aaa-server-host)# action-uri -ab41-8333b16a0008&GUID=&SMAUTHREASON
hostname(config-aaa-server-host)# action-uri =0&METHOD=GET&SMAGENTNAME=$SM$5FZmjnk
hostname(config-aaa-server-host)# action-uri 3DRNwNjk2KcqVCFbIrNT9%2bJ0H0KpshFtg6r
hostname(config-aaa-server-host)# action-uri B1UV2PxxHqLw%3d%3d&TARGET=https%3A%2F
hostname(config-aaa-server-host)# action-uri %2Fauth.example.com
hostname(config-aaa-server-host)#
```



Note

You must include the hostname and protocol in the action URI. In the preceding example, these appear at the start of the URI in `http://www.example.com`.

- Step 3** To configure a username parameter for the HTTP POST request, enter the **user-parameter** command in aaa-server-host configuration mode. For example, the following command configures the username parameter `userid`:

```
hostname(config-aaa-server-host)# user-parameter userid
hostname(config-aaa-server-host)#
```

- Step 4** To configure a user password parameter for the HTTP POST request, use the **password-parameter** command in aaa-server-host configuration mode. For example, the following command configures a user password parameter named `user_password`:

```
hostname(config-aaa-server-host)# password-parameter user_password
```

```
hostname(config-aaa-server-host)#
```

- Step 5** To specify hidden parameters for exchange with the authenticating web server, use the **hidden-parameter** command in aaa-server-host configuration mode. An example hidden parameter excerpted from a POST request follows:

```
SMENC=ISO-8859-1&SMLOCALE=US-EN&target=https%3A%2F%2Fwww.example.com%2Femco%2Fappdir%2FAreaRoot.do%3FEMCOPageCode%3DENG&smauthreason=0
```

This hidden parameter includes four form entries and their values, separated by &. The four entries and their values are:

- SMENC with a value of ISO-8859-1
- SMLOCALE with a value of US-EN
- target with a value of https%3A%2F%2Fwww.example.com%2Femco%2Fappdir%2FAreaRoot.do%3FEMCOPageCode%3DENG
- smauthreason with a value of 0

To specify this hidden parameter, enter the following commands:

```
hostname(config)# aaa-server testgrp1 host example.com
hostname(config-aaa-server-host)# hidden-parameter SMENC=ISO-8859-1&SMLOCALE=US-EN&targe
hostname(config-aaa-server-host)# hidden-parameter t=https%3A%2F%2Fwww.example.com%2Femc
hostname(config-aaa-server-host)# hidden-parameter o%2Fappdir%2FAreaRoot.do%3FEMCOPageCo
hostname(config-aaa-server-host)# hidden-parameter de%3DENG&smauthreason=0
hostname(config-aaa-server-host)#
```

- Step 6** To specify the name for the authentication cookie, enter the **auth-cookie-name** command in aaa-server-host configuration mode. This command is optional. The following example specifies the authentication cookie name of SsoAuthCookie:

```
hostname(config-aaa-server-host)# auth-cookie-name SsoAuthCookie
hostname(config-aaa-server-host)#
```

Authenticating with Digital Certificates

WebVPN users that authenticate using digital certificates do not use global authentication and authorization settings. Instead, they use an authorization server to authenticate once the certificate validation occurs. For more information on authentication and authorization using digital certificates, see [“Using Certificates and User Login Credentials”](#) in the [“Configuring AAA Servers and the Local Database”](#) chapter.

Creating and Applying WebVPN Policies

Creating and applying WebVPN policies that govern access to resources at the central site includes the following tasks:

- [Creating Port Forwarding, URL, and Access Lists in Global Configuration Mode](#)
- [Assigning Lists to Group Policies and Users in Group-Policy or User Mode](#)
- [Enabling Features for Group Policies and Users](#)
- [Assigning Users to Group Policies](#)

[Chapter 30, “Configuring Tunnel Groups, Group Policies, and Users”](#) includes step-by-step instructions for all of these tasks.

Creating Port Forwarding, URL, and Access Lists in Global Configuration Mode

Use the **port forward**, **url-list**, and **access-list** commands in global configuration mode to configure the lists of ports to forward and URLs to present to WebVPN users, and their level of access. See

Assigning Lists to Group Policies and Users in Group-Policy or User Mode

After you configure port forwarding and URL lists, use the **port forward** and **url-list**, and **filter** commands in webvpn group-policy or user mode to assign lists to group policies and/or users.

Enabling Features for Group Policies and Users

To enable features for group policies and users, issue the **functions** command in group-policy or user configuration mode.

Assigning Users to Group Policies

Assigning users to group policies simplifies the configuration by letting you apply policies to many users. You can use an internal authentication server or a RADIUS server to assign users to group policies. See [Chapter 30, “Configuring Tunnel Groups, Group Policies, and Users”](#) for a thorough explanation of ways to simplify configuration with group policies.

Using the Security Appliance Authentication Server

You can configure users to authenticate to the security appliance internal authentication server, and assign these users to a group policy on the security appliance.

Using a RADIUS Server

Using a RADIUS server to authenticate users, assign users to group policies by following these steps:

-
Step 1	Authenticate the user with RADIUS and use the Class attribute to assign that user to a particular group policy.
Step 2	Set the class attribute to the group policy name in the format <code>OU=group_name</code> For example, to set a WebVPN user to the <code>SSL_VPN</code> group, set the RADIUS Class Attribute to a value of <code>OU=SSL_VPN</code> ; (Do not omit the semicolon.)
-

Configuring WebVPN Tunnel Group Attributes

Table 37-1 provides a list of tunnel group attributes that are specific to WebVPN. In addition to these attributes, you configure general tunnel group attributes common to all VPN connections. For step-by-step information on configuring tunnel groups, see “[Configuring WebVPN Tunnel Groups](#)” in Chapter 30, “[Configuring Tunnel Groups, Group Policies, and Users](#).”

Table 37-1 WebVPN Tunnel Group Attributes

Command	Function
authentication	Sets the authentication method.
customization	Identifies the name of a previously defined customization to apply.
nbns-server	Identifies the name of the NetBIOS Name Service server (nbns-server) to use for CIFS name resolution.
group-alias	Specifies the alternate names by which the server can refer to a tunnel group
group-url	Identifies one or more group URLs. If you configure this attribute, users coming in on a specified URL need not select a group at login
dns-group	Identifies the DNS server group that specifies the DNS server name, domain name, name server, number of retries, and timeout values
hic-fail-group-policy	Specifies a VPN feature policy if you use the Cisco Secure Desktop Manager to set the Group-Based Policy attribute to “Use Failure Group-Policy” or “Use Success Group-Policy, if criteria match.”

Configuring WebVPN Group Policy and User Attributes

Table 37-2 provides a list of WebVPN group policy and user attributes. For step-by-step instructions on configuring group policy and user attributes, see “[Configuring Group Policies](#)” and “[Configuring Attributes for Specific Users](#)” in Chapter 30, “[Configuring Tunnel Groups, Group Policies, and Users](#).”

Table 37-2 WebVPN Group Policy and User Attributes

Command	Function
auto-signon	Sets values for auto signon, which requires only that a user enter username and password credentials only once for a WebVPN connection.
customization	Assigns a customization object to a group-policy or user.
deny-message	Specifies the message delivered to a remote user who logs into WebVPN successfully, but has no VPN privileges.
filter	Sets the name of the webtype access list.
functions	Enables some or all of these WebVPN features: auto-download, Citrix, file access, file browsing, file entry, filter, http-proxy, URL entry, MAPI proxy, port forwarding.
homepage	Sets the URL of the web page that displays upon login.
html-content-filter	Configures the content and objects to filter from the HTML for this group policy.
http-comp	Configures compression.

Table 37-2 WebVPN Group Policy and User Attributes

Command	Function
keep-alive-ignore	Sets the maximum object size to ignore for updating the session timer.
port-forward	Applies a list of WebVPN TCP ports to forward. The user interface displays the applications on this list.
port-forward-name	Configures the name of the port forwarding applet.
sso-server	Sets the name of the SSO server.
svc	Configures SSL VPN Client attributes.
url-list	Applies a list of WebVPN servers and URLs that the user interface displays for end user access.

Configuring Application Access

The following sections provide information about configuring application access:

[Downloading the Port-Forwarding Applet Automatically](#)

[Closing Application Access to Prevent hosts File Errors](#)

[Recovering from hosts File Errors When Using Application Access](#)

Downloading the Port-Forwarding Applet Automatically

To run a remote application over WebVPN, a user clicks **Start Application Access** on the WebVPN homepage to download and start a port-forwarding Java applet. To simplify application access and shorten start time, you can configure WebVPN to automatically download this port-forwarding applet when the user first logs in to WebVPN.

To enable automatic download of the port-forwarding applet, enter the **functions** command in webvpn mode using the **auto-download** option.

**Note**

Before you configure the auto-download feature, you must first enable an application that uses the applet: port forwarding, Outlook/Exchange proxy, or HTTP proxy.

Closing Application Access to Prevent hosts File Errors

To prevent hosts file errors that can interfere with Application Access, close the Application Access window properly when you finish using Application Access. To do so, click the close icon.

Recovering from hosts File Errors When Using Application Access

The following errors can occur if you do not close the Application Access window properly:

- The next time you try to start Application Access, it might be disabled; you receive a Backup HOSTS File Found error message.

- The applications themselves might be disabled or might malfunction, even when you are running them locally.

These errors can result from terminating the Application Access window in any improper way. For example:

- Your browser crashes while you are using Application Access.
- A power outage or system shutdown occurs while you are using Application Access.
- You minimize the Application Access window while you are working, then shut down your computer with the window active (but minimized).

This section includes the following topics:

- [Understanding the hosts File](#)
- [Stopping Application Access Improperly](#)
- [Reconfiguring a hosts File](#)

Understanding the hosts File

The hosts file on your local system maps IP addresses to host names. When you start Application Access, WebVPN modifies the hosts file, adding WebVPN-specific entries. Stopping Application Access by properly closing the Application Access window returns the file to its original state.

Before invoking Application Access...	hosts file is in original state.
When Application Access starts....	• WebVPN copies the hosts file to hosts.webvpn, thus creating a backup. • WebVPN then edits the hosts file, inserting WebVPN-specific information.
When Application Access stops...	• WebVPN copies the backup file to the <code>hosts</code> file, thus restoring the hosts file to its original state. • WebVPN deletes hosts.webvpn.
After finishing Application Access...	hosts file is in original state.



Note

Microsoft anti-spyware software blocks changes that the port forwarding JAVA applet makes to the hosts file. See www.microsoft.com for information on how to allow hosts file changes when using anti-spyware software.

Stopping Application Access Improperly

When Application Access terminates abnormally, the `hosts` file remains in a WebVPN-customized state. WebVPN checks the state the next time you start Application Access by searching for a `hosts.webvpn` file. If it finds one, a `Backup HOSTS File Found` error message ([Figure 37-4](#)) appears, and Application Access is temporarily disabled.

Once you shut down Application Access improperly, you leave your remote access client/server applications in limbo. If you try to start these applications without using WebVPN, they might malfunction. You might find that hosts that you normally connect to are unavailable. This situation could commonly occur if you run applications remotely from home, fail to quit the Application Access window before shutting down the computer, then try to run the applications later from the office.

Reconfiguring a hosts File

To reenale Application Access or malfunctioning applications:


- If you are able to connect to your remote access server, follow the steps in the section [“Reconfiguring a hosts File Automatically Using WebVPN.”](#)
- If you are unable to connect to your remote access server from your current location or if you have made custom edits to the hosts file, follow the steps in the section [“Reconfiguring hosts File Manually.”](#)

Reconfiguring a hosts File Automatically Using WebVPN

If you are able to connect to your remote access server, follow these steps to reconfigure the hosts file and reenale both Application Access and the applications.

-
- Step 1** Start WebVPN and log in. The home page opens.
- Step 2** Click the **Applications Access** link. A Backup HOSTS File Found message appears. (See [Figure 37-4](#).)

Figure 37-4 Backup HOSTS File Found Message



- Step 3** Choose one of the following options:
- **Restore from backup** — WebVPN forces a proper shutdown. WebVPN copies the hosts.webvpn backup file to the hosts file, restoring it to its original state, then deletes hosts.webvpn. You then have to restart Application Access.
 - **Do nothing** — Application Access does not start. The remote access home page reappears.
 - **Delete backup** — WebVPN deletes the hosts.webvpn file, leaving the hosts file in its WebVPN-customized state. The original hosts file settings are lost. Application Access then starts, using the WebVPN-customized hosts file as the new original. Choose this option only if you are

unconcerned about losing hosts file settings. If you or a program you use might have edited the hosts file after Application Access has shut down improperly, choose one of the other options, or edit the hosts file manually. (See the “[Reconfiguring hosts File Manually](#)” section.)

Reconfiguring hosts File Manually

If you are not able to connect to your remote access server from your current location, or if you have customized the hosts file and do not want to lose your edits, follow these steps to reconfigure the hosts file and reenab both Application Access and the applications.

Step 1 Locate and edit your hosts file. The most common location is c:\windows\system32\drivers\etc\hosts.

Step 2 Check to see if any lines contain the string: # added by WebVpnPortForward
If any lines contain this string, your hosts file is WebVPN-customized. If your hosts file is WebVPN-customized, it looks similar to the following example:

```
123.0.0.3 server1 # added by WebVpnPortForward
123.0.0.3 server1.example.com vpn3000.com # added by WebVpnPortForward
123.0.0.4 server2 # added by WebVpnPortForward
123.0.0.4 server2.example.com.vpn3000.com # added by WebVpnPortForward
123.0.0.5 server3 # added by WebVpnPortForward
123.0.0.5 server3.example.com vpn3000.com # added by WebVpnPortForward

# Copyright (c) 1993-1999 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#       102.54.94.97      cisco.example.com          # source server
#       38.25.63.10      x.example.com              # x client host

123.0.0.1      localhost
```

Step 3 Delete the lines that contain the string: # added by WebVpnPortForward

Step 4 Save and close the file.

Step 5 Start WebVPN and log in.

The home page appears.

Step 6 Click the Application Access link.

The Application Access window appears. Application Access is now enabled.

Configuring File Access

The Common Internet File System (CIFS) protocol provides users with network access to files, printers, and other machine resources. Microsoft implemented CIFS for networks of Windows computers, however, open source implementations of CIFS provide file access to servers running other operating systems, such as Linux, UNIX, and Mac OS X.

WebVPN serves remote users with HTTPS portal pages that interface with a proxy CIFS client running on the security appliance. Using this client, WebVPN provides users with network access to the files on the network, to the extent that the users meet user authentication requirements and the file properties do not restrict access. The client is transparent; the portal pages delivered by WebVPN provide the appearance of direct access to the file systems.


When a user requests a list of files, WebVPN queries the server designated as the master browser for the IP address of the server containing the list. The security appliance gets the list and delivers it to the remote user on a portal page.

WebVPN lets the user invoke the following CIFS functions, depending on user authentication requirements and file properties:

- Navigate and list domains and workgroups, servers within a domain or workgroup, shares within a server, and files within a share or directory
- Create directories
- Download, upload, rename, move, and delete files

The security appliance requires a master browser or WINS server, typically on the same network as the security appliance or reachable from that network, to query the network for a list of servers when the remote user clicks Browse Networks on the WebVPN home page or toolbar (Figure 37-5).

Figure 37-5 Browse Networks on the WebVPN Home Page and Floating Toolbar



The master browser provides the CIFS client on the security appliance with a list of the resources on the network, which WebVPN serves to the remote user. You cannot use a DNS server for a master browser. WebVPN supports file access in an Active Native Directory environment using a WINS server, but not a Dynamic DNS server.

Step 1 of the following procedure describes how to specify the master browser and WINS servers. As an alternative to following the instructions Step 1, you can use the **url-list** command in global configuration mode or in webvpn mode, which you enter from group-policy or username mode, to configure a server share in the File Folder Bookmarks. For example:

```
url-list listname displayname cifs://ServerA/ShareX/
```

Using this method (adding a share) does not require a master browser or a WINS server, however, it does not provide support for the Browse Networks link. You can use a hostname or an IP address to refer to ServerA when entering this command. If you use a hostname, the security appliance requires a DNS server to resolve it to an IP address.

**Note**

Before configuring file access, you must configure the shares on the servers for user access.

Add support for CIFS access to files as follows:

- Step 1** Use the **nbns-server** command in tunnel-group webvpn configuration mode once for each NetBIOS Name Server (NBNS).

```
nbns-server {IPaddress | hostname} [master] [timeout timeout] [retry retries]
```

master is the computer designated as the master browser. The master browser maintains the list of computers and shared resources. Any NBNS server you identify with this command without entering the master portion of the command must be a Windows Internet Naming Server (WINS). Specify the master browser first, then specify the WINS servers. You can specify up to three servers, including the master browser, for a tunnel group.

retries is the number of times to retry queries to the NBNS server. The security appliance recycles through the list of servers this number of times before sending an error message. The default value is 2; the range is 1 through 10.

timeout is the number of seconds the security appliance waits before sending the query again, to the same server if it is the only one, or another server if there are more than one. The default timeout is 2 seconds; the range is 1 to 30 seconds.

For example,

```
hostname(config-tunnel-webvpn) # nbns-server 192.168.1.20 master
hostname(config-tunnel-webvpn) # nbns-server 192.168.1.41
hostname(config-tunnel-webvpn) # nbns-server 192.168.1.47
```

**Note**

Use the **tunnel-group webvpn-attributes** command if you want to display the NBNS servers already present in the tunnel group configuration.

- Step 2** (Optional) Use the **character-encoding** command to specify the character set to encode in WebVPN portal pages to be delivered to remote users. By default, the encoding type set on the remote browser determines the character set for WebVPN portal pages, so you need to set the character encoding only if it is necessary to ensure proper encoding on the browser.

```
character-encoding charset
```

Charset is a string consisting of up to 40 characters, and equal to one of the valid character sets identified in <http://www.iana.org/assignments/character-sets>. You can use either the name or the alias of a character set listed on that page. Examples include iso-8859-1, shift_jis, and ibm850.

**Note**

The character-encoding and file-encoding values do not exclude the font family to be used by the browser. You need to complement the setting of one these values with the **page style** command in webvpn customization command mode to replace the font family if you are using Japanese Shift_JIS character encoding, as shown in the following example, or enter the **no page style** command in webvpn customization command mode to remove the font family.

The following example sets the character-encoding attribute to support Japanese Shift_JIS characters, removes the font family, and retains the default background color:

```
hostname(config-webvpn)# character-encoding shift_jis
hostname(config-webvpn)# customization DfltCustomization
hostname(config-webvpn-custom)# page style background-color:white
hostname(config-webvpn-custom)#
```

- Step 3** (Optional) Use the **file-encoding** command to specify the encoding for WebVPN portal pages from specific CIFS servers. Thus, you can use different file-encoding values for CIFS servers that require different character encodings.

file-encoding {server-name | server-ip-address} charset

The following example sets the file-encoding attribute of the CIFS server 10.86.5.174 to support IBM860 (alias “CP860”) characters:

```
hostname(config-webvpn)# file-encoding 10.86.5.174 cp860
hostname(config-webvpn)#
```

- Step 4** To configure security appliance support for file access, file browsing, and file server entry, use the **functions** command in webvpn mode, which you enter from group-policy or username mode.

functions file-access file-browsing file-entry

For example:

```
hostname(config-group-webvpn)# functions file-access file-browsing file-entry
hostname(config-group-policy)#
```

For a complete description of these commands, see the *Cisco Security Appliance Command Reference*.

Configuring Access to Citrix MetaFrame Services

WebVPN users can use a connection to the security appliance to access Citrix MetaFrame services. In this configuration, the security appliance functions as the Citrix secure gateway. Complete the following steps to configure support for Citrix MetaFrame services:

- Step 1** Configure the Citrix Web Interface software to operate in a mode that does not use the secure gateway.
- Step 2** Install an SSL certificate onto the security appliance interface to which remote users use a fully-qualified domain name (FQDN) to connect.

**Note**

Do not specify an IP address as the common name (CN) for the SSL certificate. The remote user attempts to use the FQDN to communicate with the security appliance. The remote PC must be able to use DNS or an entry in the System32\drivers\etc\hosts file to resolve the FQDN.

- Step 3** Use the **functions citrix** command once for each group policy or user for which you want to enable Citrix support.

The following example shows how to configure Citrix for a group policy named FirstGroup:

```
hostname(config)# group-policy FirstGroup internal
hostname(config)# group-policy FirstGroup attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# functions citrix
hostname(config-group-webvpn)#
```

Using WebVPN with PDAs

You can access WebVPN from your Pocket PC or other certified personal digital assistant device. Neither the security appliance administrator nor the WebVPN user need do anything special to use WebVPN with a certified PDA.

Cisco has certified the following PDA platform:

```
HP iPaq H4150
Pocket PC 2003
Windows CE 4.20.0, build 14053
Pocket Internet Explorer (PIE)
ROM version 1.10.03ENG
ROM Date: 7/16/2004
```

Some differences in the PDA version of WebVPN exist:

- A banner web page replaces the popup WebVPN window.
- An icon bar replaces the standard WebVPN floating toolbar. This bar displays the Go, Home and Logout buttons.
- The Show Toolbar icon is not included on the main WebVPN portal page.
- Upon WebVPN logout, a warning message provides instructions for closing the PIE browser properly. If you do not follow these instructions and you close the browser window in the common way, PIE does not disconnect from WebVPN or any secure website that uses HTTPS.
- WebVPN supports OWA 2000 and OWA 2003 Basic Authentication. If Basic Authentication is not configured on an OWA server and a WebVPN user attempts to access that server, access is denied.
- Unsupported WebVPN features:
 - Application Access (port forwarding) and other Java-dependent features
 - MAPI proxy
 - HTTP proxy
 - Cisco Secure Desktop (CSD does provide limited support for Microsoft Windows CE)
 - Microsoft Outlook Web Access (OWA) 5.5

- The Citrix Metaframe feature (if the PDA does not have the corresponding Citrix ICA client software)

Using E-Mail over WebVPN

WebVPN supports several ways to access e-mail. This section includes the following methods:

- [Configuring E-mail Proxies](#)
- [Configuring MAPI](#)
- [Configuring Web E-mail: MS Outlook Web Access](#)

Configuring E-mail Proxies

WebVPN supports IMAP4S, POP3S, and SMTPS e-mail proxies. [Table 37-3](#) lists attributes that apply globally to e-mail proxy users:

Table 37-3 WebVPN Attributes for E-mail Proxy Users

Function	Command	Default Value
Specifies the previously configured accounting servers to use with e-mail proxy.	accounting-server-group	None
Specifies the authentication method(s) for e-mail proxy users.	authentication	IMAP4S: Mailhost (required) POP3S Mailhost (required) SMTPS: AAA
Specifies the previously configured authentication servers to use with e-mail proxy.	authentication-server-group	LOCAL
Specifies the previously configured authorization servers to use with WebVPN.	authorization-server-group	None
Requires users to authorize successfully to connect.	authorization-required	Disabled
Identifies the DN of the peer certificate to use as a username for authorization.	authorization-dn-attributes	Primary attribute: CN Secondary attribute: OU
Specifies the name of the group policy to use.	default-group-policy	DfltGrpPolicy
Enables e-mail proxy on the specified interface.	enable	Disabled
Defines the separator between the e-mail and VPN usernames and passwords.	name-separator	“:” (colon)
Configures the maximum number of outstanding non-authenticated sessions.	outstanding	20
Sets the port the e-mail proxy listens to.	port	IMAP4S:993 POP3S: 995 SMTPS: 988 ¹

Table 37-3 WebVPN Attributes for E-mail Proxy Users

Function	Command	Default Value
Specifies the default e-mail server.	server	None.
Defines the separator between the e-mail and server names.	server-separator	“@”

1. With the Eudora e-mail client, SMTPS works only on port 465, even though the default port for SMTPS connections is 988.

E-mail Proxy Certificate Authentication

Certificate authentication for e-mail proxy connections works with Netscape 7x e-mail clients. Other e-mail clients such as MS Outlook, MS Outlook Express, and Eudora lack the ability to access the certificate store.

Configuring MAPI

MAPI, also called MS Outlook Exchange proxy, has the following requirements:

- MS Outlook Exchange must be installed on the remote computer.
- You must enable MS Outlook Exchange Proxy on a security appliance interface. You do this by entering the **functions** command, which is a group-policy webvpn command. For example:

```
hostname(config)# group-policy group_policy_name attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# functions mapi
```

- Provide the Exchange server NetBIOS name. The Exchange server must be on the same domain as the security appliance DNS server. For example:

```
hostname(config)# domain domain_name
hostname(config)#
```



Note

An open MS Outlook client connected via MS Outlook Exchange Mail Proxy checks continually for mail on the Exchange Server, which keeps the connection open. As long as Outlook is open, the connection never times out, regardless of the settings.

Configuring Web E-mail: MS Outlook Web Access

Web e-mail is MS Outlook Web Access for Exchange 2000, Exchange 5.5, and Exchange 2003. It requires an MS Outlook Exchange Server at the central site. It also requires that users perform the following tasks:

- Enter the URL of the mail server in a browser in your WebVPN session.
- When prompted, enter the e-mail server username in the format *domain\username*.
- Enter the e-mail password.

Optimizing WebVPN Performance

The security appliance provides several ways to optimize WebVPN performance and functionality. Performance improvements include caching and compressing web objects. Functionality tuning includes setting limits on content transformation and proxy-bypass. APCF provides an additional method of tuning content transformation. The following sections explain these features:

- [Configuring Caching](#)
- [Configuring Content Transformation](#)

Configuring Caching

Caching enhances WebVPN performance. It stores frequently reused objects in the system cache, which reduces the need to perform repeated rewriting and compressing of content. It reduces traffic between WebVPN and the remote servers, with the result that many applications run much more efficiently.

By default, caching is enabled. You can customize the way caching works for your environment by using the caching commands in cache mode, which you enter from webvpn mode, as in the following example.

```
hostname(config)#  
hostname(config)# webvpn  
hostname(config-webvpn)# cache  
hostname(config-webvpn-cache)#
```

A list of caching commands and their functions follows:

Cache Command	Function
cache-compressed	Caches compressed content.
disable	Disables caching.
expiry-time	Configures an expiration time for caching objects.
lmfactor	Configures terms for revalidating cached objects.
max-object-size	Sets a maximum size for objects to cache.
min-object-size	Sets a minimum size for objects to cache.

Configuring Content Transformation

By default, the security appliance processes all WebVPN traffic through a content transformation/rewriting engine that includes advanced elements such as JavaScript and Java to proxy HTTP traffic that may have different semantics and access control rules depending on whether the user is accessing an application within or independently of an SSL VPN device.

Some web resources require highly individualized treatment. The following sections describe functionality that provides such treatment:

- [Configuring a Certificate for Signing Rewritten Java Content](#)
- [Disabling Content Rewrite](#)
- [Using Proxy Bypass](#)
- [Configuring Application Profile Customization Framework](#)

Subject to the requirements of your organization and the web content involved, you might use one of these features.

Configuring a Certificate for Signing Rewritten Java Content

Java objects which have been transformed by WebVPN can subsequently be signed using a PKCS12 digital certificate associated with a trustpoint. You import and employ the certificate using a combination of the **crypto ca import** and **java-trustpoint** commands.

The following example commands show the creation of a trustpoint named **mytrustpoint** and its assignment to signing Java objects:

```
hostname(config)# crypto ca import mytrustpoint pkcs12 mypassphrase
Enter the base 64 encoded PKCS12.
End with the word "quit" on a line by itself.
[ PKCS12 data omitted ]
quit
INFO: Import PKCS12 operation completed successfully.
hostname(config)# webvpn
hostname(config)# java-trustpoint mytrustpoint
hostname(config)#
```

Disabling Content Rewrite

You might not want some applications and web resources, for example, public websites, to go through the security appliance. The security appliance therefore lets you create rewrite rules that let users browse certain sites and applications without going through the security appliance. This is similar to split-tunneling in an IPsec VPN connection.

Use the **rewrite** command with the **disable** option in **webvpn** mode to specify applications and resources to access outside a WebVPN tunnel.

You can use the **rewrite** command multiple times. The order number of rules is important because the security appliance searches rewrite rules by order number, starting with the lowest, and applies the first rule that matches.

Using Proxy Bypass

You can configure the security appliance to use proxy bypass when applications and web resources work better with the special content rewriting this feature provides. Proxy bypass is an alternative method of content rewriting that makes minimal changes to the original content. It is often useful with custom web applications.

You can use this command multiple times. The order in which you configure entries is unimportant. The interface and path mask or interface and port uniquely identify a proxy bypass rule.

If you configure proxy bypass using ports rather than path masks, depending on your network configuration, you might need to change your firewall configuration to allow these ports access to the security appliance. Use path masks to avoid this restriction. Be aware, however, that path masks can change, so you might need to use multiple pathmask statements to exhaust the possibilities.

A path is everything in a URL after the .com or .org or other types of domain name. For example, in the URL **www.mycompany.com/hrbenefits**, **hrbenefits** is the path. Similarly, for the URL **www.mycompany.com/hrinsurance**, **hrinsurance** is the path. If you want to use proxy bypass for all hr sites, you can avoid using the command multiple times by using the * wildcard as follows: **/hr***.

To configure proxy bypass, use the **proxy-bypass** command in **webvpn** mode.

Configuring Application Profile Customization Framework

An APCF profile for WebVPN lets the security appliance handle non-standard applications and web resources so that they display correctly over a WebVPN connection. An APCF profile contains a script that specifies when (pre, post), where (header, body, request, response), and what data to transform for a particular application. The script is in XML and uses sed (stream editor) syntax for string/text transformation. Multiple APCF profiles can run in parallel on a security appliance. Within an APCF profile script, multiple APCF rules can apply. In this case, the security appliance processes the oldest rule first (based on configuration history), then the next oldest rule, and so forth.

You can store APCF profiles on the security appliance flash memory, or on an HTTP, HTTPS, or TFTP server. Use the **apcf** command in webvpn mode to identify and locate an APCF profile that you want to load on the security appliance.

The following example shows how to enable an APCF profile named apcf1.xml, located on flash memory.

```
hostname(config)# webvpn
hostname(config-webvpn)# apcf flash:/apcf/apcf1.xml
hostname(config-webvpn)#
```

This example shows how to enable an APCF profile named apcf2.xml, located on an https server called myserver, port 1440 with the path being /apcf.

```
hostname(config)# webvpn
hostname(config-webvpn)# apcf https://myserver:1440/apcf/apcf2.xml
hostname(config-webvpn)#
```

APCF Syntax

The following sections describe APCF syntax.



Caution

Misuse of an APCF profile can result in reduced performance and undesired rendering of content. In most cases, Cisco Engineering supplies APCF profiles to solve specific application rendering issues.

APCF profiles use XML format, and sed script syntax, with the XML tags in [Table 37-4](#)

Table 37-4 **APCF XML Tags**

Tag	Use
<APCF>...</APCF>	The mandatory root element that opens any APCF XML file.
<version>1.0</version>	The mandatory tag that specifies the APCF implementation version. Currently the only version is 1.0.
<application>...</application>	The mandatory tag that wraps the body of the XML description.
<id> text </id>	The mandatory tag that describes this particular APCF functionality.
<apcf-entities>...</apcf-entities>	The mandatory tag that wraps a single or multiple APCF entities.

Table 37-4 **APCF XML Tags (continued)**

Tag	Use
<code><js-object>...</js-object></code> <code><html-object>...</html-object></code> <code><process-request-header>...</process-request-header></code> <code><process-response-header>...</process-response-header></code> <code><preprocess-request-body>...</preprocess-request-body></code> <code><postprocess-request-body>...</postprocess-request-body></code> <code><preprocess-response-body>...</preprocess-response-body></code> <code><postprocess-response-body>...</postprocess-response-body></code>	One of these tags specifying type of content or the stage at which the APCF processing should take place is required.
<code><conditions>... </conditions></code>	A child element of the pre/post-process tags that specifies criteria for processing such as: http-version (such as 1.1, 1.0, 0.9) http-method (get, put, post, webdav) http-scheme (http, https, other) server-regexp regular expression containing ("a".. "z" "A".. "Z" "0".. "9" ".-_*[]?")) server-fnmatch (regular expression containing ("a".. "z" "A".. "Z" "0".. "9" ".-_*[]?+()\{\},"), user-agent-regexp user-agent-fnmatch request-uri-regexp request-uri-fnmatch If more than one of condition tags are present, the security appliance performs a logical AND for all tags.
<code><action> ... </action></code>	Wraps one or more actions to perform on the content under specified conditions; define each of these actions with the following <code><do></code> tag or the <code><sed-script></code> tag.
<code><do>...</do></code>	Defines one of the following actions: <code><no-rewrite></code> <code><no-toolbar></code> <code><no-gzip></code> <code><force-cache></code> <code><force-no-cache></code>
<code><sed-script> TEXT </sed-script></code>	The child element of the action tag. The TEXT must be a valid Sed script. The <code><sed-script></code> applies to the <code><conditions></code> tag defined before it.

APCF Example

An example of an APCF profile follows:

```
<APCF>
<version>1.0</version>
<application>
  <id>Do not compress content from notsogood.com</id>
  <apcf-entities>
    <process-request-header>
      <conditions>
        <server-fnmatch>*.notsogood.com</server-fnmatch>
      </conditions>
      <action>
        <do><no-gzip/></do>
      </action>
    </process-request-header>
  </apcf-entities>
</application>
</APCF>
```

WebVPN End User Setup

This section is for the system administrator who sets up WebVPN for end users. It describes how to customize the end-user interface.

This section summarizes configuration requirements and tasks for a remote system. It specifies information to communicate to users to get them started using WebVPN. It includes the following topics:

- [Defining the End User Interface](#)
- [Customizing WebVPN Pages, page 37-35](#)
- [Requiring Usernames and Passwords](#)
- [Communicating Security Tips](#)
- [Configuring Remote Systems to Use WebVPN Features](#)

Defining the End User Interface

The WebVPN end user interface consists of a series of HTML panels. A user logs on to WebVPN by entering the IP address of a security appliance interface in the format `https://address`. The first panel that displays is the login screen ([Figure 37-6](#)).

Figure 37-6 WebVPN Login Screen




The screenshot shows a web browser window with a title bar that says "Login". Inside the window, there is a message "Please enter your username and password." followed by three input fields labeled "USERNAME:", "PASSWORD:", and "GROUP:". Below these fields are two buttons: "Login" and "Clear".

153013

Viewing the WebVPN Home Page

After the user logs in, the WebVPN Home page opens (Figure 37-7).


Figure 37-7 WebVPN Home Page



The home page displays all of the WebVPN features you have configured, and its appearance reflects the logo, text, and colors you have selected. This sample home page includes all available WebVPN features with the exception of identifying specific file shares. It lets users browse the network, enter URLs, access specific websites, and use port forwarding to access TCP applications.

Viewing the WebVPN Application Access Panel

To start port forwarding, also called application access, a user clicks the Go button in the Application Access box. The Application Access window opens (Figure 37-8).

Figure 37-8 WebVPN Application Access Window

This window displays the TCP applications configured for this WebVPN connection. To use an application with this panel open, the user starts the application in the normal way.

Viewing the Floating Toolbar

The floating toolbar shown in Figure 37-9 represents the current WebVPN session.

Figure 37-9 WebVPN Floating Toolbar

Be aware of the following characteristics of the floating toolbar:

- The toolbar lets you enter URLs, browse file locations, and choose preconfigured web connections without interfering with the main browser window.
- If you configure your browser to block popups, the floating toolbar cannot display.
- If you close the toolbar, the security appliance prompts you to confirm that you want to end the WebVPN session.

See [Table 37-6 on page 37-46](#) for detailed information about using WebVPN.

Customizing WebVPN Pages

You can change the appearance of WebVPN pages displayed to WebVPN users. This includes the Login page displayed to users when they connect to the security appliance, the Home page displayed to users after the security appliance authenticates them, the Application Access window displayed when users launch an application, and the Logout page displayed when users logout of WebVPN service.

After you customize the WebVPN pages, you can save your customization and apply it to a specific tunnel group, group, or user. You can create and save many customizations, enabling the security appliance to change the appearance of WebVPN pages for individual users, or group of users.

This section contains the following topics and tasks:

- [Using Cascading Style Sheet Parameters, page 37-35](#)
- [Customizing the WebVPN Login Page, page 37-36](#)
- [Customizing the WebVPN Logout Page, page 37-37](#)
- [Customizing the WebVPN Home Page, page 37-38](#)
- [Customizing the Application Access Window, page 37-40](#)
- [Customizing the Prompt Dialogs, page 37-41](#)
- [Applying Customizations to Tunnel Groups, Groups and Users, page 37-42](#)

Using Cascading Style Sheet Parameters

Many WebVPN customization commands contain the **style** option. The value is expressed as any valid Cascading Style Sheet (CSS) parameters. Describing these parameters is beyond the scope of this document. For more information about CSS parameters, consult CSS specifications at the World Wide Web Consortium (W3C) website at www.w3.org. Appendix F of the CSS 2.1 Specification contains a convenient list of CSS parameters, and is available at www.w3.org/TR/CSS21/propidx.html.

Here are some tips for making the most common changes to the WebVPN pages—the page colors:

- You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.
- RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.
- HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.




Note

To easily customize the WebVPN pages, we recommend that you use ASDM, which has convenient features for configuring style elements, including color swatches and preview capabilities.

Customizing the WebVPN Login Page

Figure 37-10 shows the WebVPN Login page and associated CLI commands that you can use to customize the page.

Figure 37-10 WebVPN Login Page and Associated CLI Commands



The following procedure guides you through customizing every element of the WebVPN Login page using CLI commands and includes examples of the commands:

Step 1 Enter WebVPN customization mode using the **customization** command from webvpn mode:

```
hostname(config)# webvpn
hostname(config-webvpn)# customization cisco
hostname(config-webvpn-custom)#
```

Step 2 Change the CSS style of the WebVPN Login page using the **page style** command:

[no] page style value

```
hostname(config-webvpn-custom)# page style font-size:large
```

Step 3 Change the title using the **title** command:

[no] title {text | style} value

```
hostname(config-webvpn-custom)# title text Cisco WebVPN Service
```

Step 4 Change the logo with a logo residing in flash memory using the **logo** command:

[no] logo {none | file {path value}}

To disallow a logo and prevent inheriting a logo, use the **none** option to set a null value.

```
hostname(config-webvpn-custom)# logo file disk0:cisco_logo.gif
```

Step 5 Change the title of the Login box using the **login-title** command:

[no] login-title {text | style} value


```
hostname(config-webvpn-custom)# login-title style background-color: rgb(51,51,255);color:
rgb(51,51,255); font-family: Algerian; font-size: 12pt; font-style: italic; font-weight:
bold
```

Step 6 Change the message of the Login box using the **login-message** command:

```
[no] login-message {text | style} value
```

```
hostname(config-webvpn-custom)# login-message text username and password
```

Step 7 Change the username prompt of the Login box using the **username-prompt** command:

```
[no] username-prompt {text | style} value
```

```
hostname(config-webvpn-custom)# username-prompt text Corporate Username:
```

```
hostname(config-webvpn-custom)# username-prompt style font-weight:bolder
```

Step 8 Change the password prompt of the Login box using the **password-prompt** command:

```
[no] password-prompt {text | style} value
```

```
hostname(config-webvpn-custom)# password-prompt text Corporate Username:
```

```
hostname(config-webvpn-custom)# password-prompt style font-weight:bolder
```

Step 9 Change the group prompt of the Login box using the **group-prompt** command:

```
[no] group-prompt {text | style} value
```

```
hostname(config-webvpn-custom)# group-prompt text Corporate Group:
```

```
hostname(config-webvpn-custom)# group-prompt style font-weight:bolder
```

Step 10 Change the content or appearance of the Login button of the Login box using the **login-button** command:

```
[no] login-button {text | style} value
```

```
hostname(config-webvpn-custom)# login-button text OK
```


Step 11 Change the content or appearance of the Clear button of the Login box using the **clear-button** command:

```
[no] clear-button {text | style} value
```

```
hostname(config-webvpn-custom)# clear-button background-color:blue
```

Customizing the WebVPN Logout Page

The security appliance displays the WebVPN Logout page when WebVPN users log out of WebVPN service. [Figure 37-11](#) shows the WebVPN Logout page and the associated CLI commands that you can use to customize the page.

Figure 37-11 WebVPN Logout Page

The following procedure guides you through customizing the WebVPN Logout page using CLI commands and includes examples of the commands:

Step 1 Enter WebVPN customization mode using the **customization** command from webvpn mode:

```
hostname(config)# webvpn
hostname(config-webvpn)# customization cisco
hostname(config-webvpn-custom)#
```

Step 2 Change the title of the Logout box using the **logout-title** command:

[no] **logout-title {text | style} value**

```
hostname(config-webvpn-custom)# logout-title style background-color: rgb(51,51,255);color:
rgb(51,51,255); font-family: Algerian; font-size: 12pt; font-style: italic; font-weight:
bold
```


Step 3 Change the message of the Logout box using the **logout-message** command:

[no] **logout-message {text | style} value**

```
hostname(config-webvpn-custom)# login-title style background-color: rgb(51,51,255);color:
rgb(51,51,255); font-family: Algerian; font-size: 12pt; font-style: italic; font-weight:
bold
```

Customizing the WebVPN Home Page

You can customize the appearance of the WebVPN Home page that the security appliance displays to authenticated WebVPN users. Figure 37-12 shows the WebVPN Home page and associated CLI commands that you can use to customize the page.

Figure 37-12 WebVPN Home Page and Associated CLI Commands

The following procedure guides you through customizing every element of the WebVPN Home page using CLI commands and includes examples of the commands:

Step 1 Enter WebVPN customization mode using the **customization** command from webvpn mode:

```
hostname(config)# webvpn
hostname(config-webvpn)# customization cisco
hostname(config-webvpn-custom)#
```

Step 2 Change the border of the WebVPN page using the **border style** command and CSS parameters:

[no] border style value

```
hostname(config-webvpn-custom)# border style background-color:66FFFF
```

Step 3 Change the appearance of the Web Applications box using the **web-applications** command:

[no] web-applications {title | message | dropdown} {text | style} value

```
hostname(config-webvpn-custom)# web-applications title text WWW Applications
hostname(config-webvpn-custom)# web-applications title style color:blue
hostname(config-webvpn-custom)# web-applications message text Enter URL
hostname(config-webvpn-custom)# web-applications message style color:blue
hostname(config-webvpn-custom)# web-applications dropdown text URLs to Browse
hostname(config-webvpn-custom)# web-applications dropdown style color:red
```

Step 4 Change the appearance of the Application Access box using the **application-access** command:

[no] application-access {title | message} {text | style} value

```
hostname(config-webvpn-custom)# application-access title text Applications
hostname(config-webvpn-custom)# application-access title style color:blue
hostname(config-webvpn-custom)# application-access message text Start Application
```

148905

```
hostname(config-webvpn-custom)# application-access message style color:blue
```

Step 5 Change the appearance of the Browse Networks box using the **browse-networks** command:

[no] browse-networks {title | message | dropdown} {text | style} value

```
hostname(config-webvpn-custom)# browse-networks title text Corporate Nets
hostname(config-webvpn-custom)# browse-networks title style color:blue
hostname(config-webvpn-custom)# browse-networks message text Enter URL
hostname(config-webvpn-custom)# browse-networks message style color:blue
hostname(config-webvpn-custom)# browse-networks dropdown text URLs to Browse
hostname(config-webvpn-custom)# browse-networks dropdown style color:red
```

Step 6 Change the Web Bookmarks title or links using the **web-bookmarks** command:

[no] web-bookmarks {link {style value} | title {style value | text value}}

```
hostname(config-webvpn-custom)# web-bookmarks link style color:black
hostname(config-webvpn-custom)# web-bookmarks title style color:black
hostname(config-webvpn-custom)# web-bookmarks title text Corporate Web Bookmarks
```

Step 7 Change the File Bookmarks title or the File Bookmarks links using the **file-bookmarks** command:

[no] file-bookmarks {link {style value} | title {style value | text value}}


```
hostname(config-webvpn-custom)# file-bookmarks link style color:blue
hostname(config-webvpn-custom)# file-bookmarks title style color:blue
hostname(config-webvpn-custom)# file-bookmarks title text Corporate File Bookmarks
```

Customizing the Application Access Window

You can customize the Application Access window that launches when the remote user selects an application. [Figure 37-13](#) shows the Application Access window and the associated CLI commands that you can use to customize it.

Figure 37-13 Application Access Window

application-access window {text | style} value



The following procedure guides you through customizing the Application Access window using CLI commands and includes examples of the commands:

Step 1 Enter WebVPN customization mode using the **customization** command from webvpn mode:

```
hostname(config)# webvpn
hostname(config-webvpn)# customization cisco
hostname(config-webvpn-custom)#
```

Step 2 Change the Application Access window using the **application-access window** command:

[no] application-access window {text | style} value

```
hostname(config-webvpn-custom)# application-access window text URLs to Browse
hostname(config-webvpn-custom)# application-access window style color:red
```

Step 3 Enable or disable the Hiding of application details displayed in the WebVPN Applications Access window using the **application-access hide-details** command:


[no] application-access hide-details {enable | disable}

The default is disabled. Application details are not hidden—they display in the Application Access window.

```
hostname(config-webvpn-custom)# application-access hide-details enable
```

Customizing the Prompt Dialogs

The security appliance may send WebVPN users various prompt dialog messages as notices or warnings. [Figure 37-14](#) shows a sample dialog message and the associated CLI commands you can use to customize the appearance of these messages.

Figure 37-14 Dialog Message and Associated CLI Command

The following procedure customizes every element of the dialog message and includes examples of the commands:

- Step 1** Enter WebVPN customization mode using the **customization** command from webvpn mode:

```
hostname(config)# webvpn
hostname(config-webvpn)# customization cisco
hostname(config-webvpn-custom)#
```

- Step 2** Customize the border of the dialog messages with the **dialog border** command:

[no] **dialog border style value**

```
hostname(config-webvpn-custom)# dialog border style color:blue
```

- Step 3** Change the appearance of the title using the **dialog title** command:

[no] **dialog title style value**

```
hostname(config-webvpn-custom)# dialog title style font:bolder
```

- Step 4** Change the appearance of the message using the **dialog message** command:

[no] **dialog message style value**

```
hostname(config-webvpn-custom)# dialog message style font:italic
```

Applying Customizations to Tunnel Groups, Groups and Users

After you create a customization, you can apply the customization to a tunnel group, a group, or a user, with the **customization** command. The options displayed with this command are different depending on the mode you are in.

For more information about configuring tunnel groups, group policies, and users, see [Chapter 30, “Configuring Tunnel Groups, Group Policies, and Users”](#).

Applying Customizations to Tunnel Groups

To apply a customization to a tunnel group, use the **customization** command from tunnel group webvpn mode:

[no] **customization name**

name is the name of a customization to apply to the tunnel group.

To remove the command from the configuration, and remove a customization from the tunnel group, use the **no** form of the command.

Enter the **customization** command followed by a question mark (?) to view a list of existing customizations.

In the following example, the user enters tunnel group webvpn mode and enables the customization *cisco* for the tunnel group *cisco_telecommuters*:

```
hostname(config)# tunnel-group cisco_telecommuters webvpn-attributes
hostname(tunnel-group-webvpn)# customization cisco
```

Applying Customizations to Groups and Users

To apply a customization to a group or user, use the **customization** command from group policy webvpn mode or username webvpn mode. In these modes, the **none** and **value** options are included:

[no] customization {none | value name}

none disables the customization for the group or user, prevents the value from being inherited, and displays the default WebVPN pages.

value name is the name of a customization to apply to the group or user.

To remove the command from the configuration, and cause the value to be inherited, use the **no** form of the command.

Enter the **customization value** command followed by a question mark (?) to view a list of existing customizations.

In the following example, the user enters group policy webvpn mode, queries the security appliance for a list of customizations, and enables the customization *cisco* for the group policy *cisco_sales*:

```
hostname(config)# group-policy cisco_sales attributes
hostname(config-group-policy)# webvpn
hostname(config-username-webvpn)# customization value ?

config-username-webvpn mode commands/options:
Available configured customization profiles:
  DfltCustomization
  cisco
hostname(config-group-webvpn)# customization value cisco
```

In the next example, the user enters username webvpn mode and enables the customization *cisco* for the user *cisco_employee*:

```
hostname(config)# username cisco_employee attributes
hostname(config-username)# webvpn
hostname(config-username-webvpn)# customization value cisco
```

Requiring Usernames and Passwords

Depending on your network, during a remote session users might have to log in to any or all of the following: the computer itself, an Internet service provider, WebVPN, mail or file servers, or corporate applications. Users might have to authenticate in many different contexts, requiring different information, such as a unique username, password, or PIN.

Table 37-5 lists the type of usernames and passwords that WebVPN users might need to know.

Table 37-5 *Username and Passwords to Give to WebVPN Users*

Login Username/ Password Type	Purpose	Entered When
Computer	Access the computer	Starting the computer
Internet Service Provider	Access the Internet	Connecting to an Internet service provider
WebVPN	Access remote network	Starting WebVPN
File Server	Access remote file server	Using the WebVPN file browsing feature to access a remote file server
Corporate Application Login	Access firewall-protected internal server	Using the WebVPN web browsing feature to access an internal protected website
Mail Server	Access remote mail server via WebVPN	Sending or receiving e-mail messages

Communicating Security Tips

Advise users always to log out from the WebVPN session. (To log out of WebVPN, click the logout icon on the WebVPN toolbar or close the browser.)

Advise users that using WebVPN does not ensure that communication with every site is secure. WebVPN ensures the security of data transmission between the remote PC or workstation and the security appliance on the corporate network. If a user then accesses a non-HTTPS web resource (located on the Internet or on the internal network), the communication from the corporate security appliance to the destination web server is not secure.

Configuring Remote Systems to Use WebVPN Features

Table 37-6 includes the following information about setting up remote systems to use WebVPN:

- Starting WebVPN
- Using the WebVPN Floating Toolbar
- Web Browsing
- Network Browsing and File Management
- Using Applications (Port Forwarding)
- Using E-mail via Port Forwarding
- Using E-mail via Web Access
- Using E-mail via e-mail proxy

Table 37-6 also provides information about the following:

- WebVPN requirements, by feature
- WebVPN supported applications
- Client application installation and configuration requirements
- Information you might need to provide end users
- Tips and use suggestions for end users

It is possible you have configured user accounts differently and that different WebVPN features are available to each user. [Table 37-6](#) organizes information by feature, so you can skip over the information for unavailable features.

Table 37-6 WebVPN Remote System Configuration and End User Requirements

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Starting WebVPN	Connection to the Internet	Any Internet connection is supported, including: • Home DSL, cable, or dial-ups • Public kiosks • Hotel hook-ups • Airport wireless nodes • Internet cafes
	WebVPN-supported browser	We recommend the following browsers for WebVPN. Other browsers might not fully support WebVPN features. On Microsoft Windows: • Internet Explorer version 6.0 • Netscape version 7.2 • Mozilla version 1.7 and later • Firefox 1.x On Linux: • Mozilla version 1.7 • Netscape version 7.2 • Firefox 1.x On Solaris: • Netscape version 7.2 On Macintosh OS X: • Safari version 1.0 • Firefox 1.x
	Cookies enabled on browser	Cookies must be enabled on the browser in order to access applications via port forwarding.
	URL for WebVPN	An https address in the following form: <code>https://address</code> where <i>address</i> is the IP address or DNS hostname of an interface of the security appliance (or load balancing cluster) on which WebVPN is enabled. For example: <code>https://10.89.192.163</code> or <code>https://cisco.example.com</code> .
	WebVPN username and password	
	[Optional] Local printer	WebVPN does not support printing from a web browser to a network printer. Printing to a local printer is supported.

Table 37-6 WebVPN Remote System Configuration and End User Requirements (continued)

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Using the WebVPN Floating Toolbar		A floating toolbar is available to simplify the use of WebVPN. The toolbar lets you enter URLs, browse file locations, and choose preconfigured web connections without interfering with the main browser window. If you configure your browser to block popups, the floating toolbar cannot display. The floating toolbar represents the current WebVPN session. If you click the Close button, the security appliance prompts you to confirm that you want to close the WebVPN session.  Tip TIP: To paste text into a text field, use Ctrl-V. (Right-clicking is disabled on the WebVPN toolbar.)
Web Browsing	Username and passwords for protected websites	Using WebVPN does not ensure that communication with every site is secure. See “Communicating Security Tips.” The look and feel of web browsing with WebVPN might be different from what users are accustomed to. For example: • The WebVPN title bar appears above each web page • You access websites by: – Entering the URL in the Enter Web Address field on the WebVPN Home page – Clicking on a preconfigured website link on the WebVPN Home page – Clicking a link on a webpage accessed via one of the previous two methods Also, depending on how you configured a particular account, it might be that: • Some websites are blocked • Only the websites that appear as links on the WebVPN Home page are available

Table 37-6 *WebVPN Remote System Configuration and End User Requirements (continued)*

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Network Browsing and File Management	File permissions configured for shared remote access	Only shared folders and files are accessible via WebVPN.
	Server name and passwords for protected file servers	—
	Domain, workgroup, and server names where folders and files reside	Users might not be familiar with how to locate their files through your organization network.
	—	Do not interrupt the Copy File to Server command or navigate to a different screen while the copying is in progress. Interrupting the operation can cause an incomplete file to be saved on the server.

Table 37-6 WebVPN Remote System Configuration and End User Requirements (continued)

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Using Applications (called Port Forwarding or Application Access)	Note On Macintosh OS X, only the Safari browser supports this feature.	
	Note Because this feature requires installing Sun Microsystems Java™ Runtime Environment and configuring the local clients, and because doing so requires administrator permissions on the local system, it is unlikely that users will be able to use applications when they connect from public remote systems.	
	 Caution Users should always close the Application Access window when they finish using applications by clicking the Close icon. Failure to quit the window properly can cause Application Access or the applications themselves to be disabled. See Recovering from hosts File Errors When Using Application Access for details.	
	Client applications installed	—
	Cookies enabled on browser	—
	Administrator privileges	User must have administrator access on the PC if you use DNS names to specify servers because modifying the hosts file requires it.
	Sun Microsystems Java Runtime Environment (JRE) version 1.4.x and 1.5.x installed. Javascript must be enabled on the browser. By default, it is enabled.	If JRE is not installed, a pop-up window displays, directing users to a site where it is available. On rare occasions, the WebVPN port forwarding applet fails with JAVA exception errors. If this happens, do the following: 1. Clear the browser cache and close the browser. 2. Verify that no JAVA icons are in the computer task bar. Close all instances of JAVA. 3. Establish a WebVPN session and launch the port forwarding JAVA applet.
	Client applications configured, if necessary. Note The Microsoft Outlook client does not require this configuration step. All non-Windows client applications require configuration. To see if configuration is necessary for a Windows application, check the value of the Remote Server. • If the Remote Server contains the server hostname, you do not need to configure the client application. • If the Remote Server field contains an IP address, you must configure the client application.	To configure the client application, use the server's locally mapped IP address and port number. To find this information: 1. Start WebVPN on the remote system and click the Application Access link on the WebVPN Home page. The Application Access window appears. 2. In the Name column, find the name of the server you want to use, then identify its corresponding client IP address and port number (in the Local column). 3. Use this IP address and port number to configure the client application. Configuration steps vary for each client application.
	Note Clicking a URL (such as one in an -e-mail message) in an application running over WebVPN does not open the site over WebVPN. To open a site over WebVPN, cut and paste the URL into the Enter WebVPN (URL) Address field.	

Table 37-6 *WebVPN Remote System Configuration and End User Requirements (continued)*

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Using E-mail via Application Access	Fulfill requirements for Application Access (See Using Applications)	To use mail, start Application Access from the WebVPN Home page. The mail client is then available for use.
	Note If you are using an IMAP client and you lose your mail server connection or are unable to make a new connection, close the IMAP application and restart WebVPN. Other mail clients	We have tested Microsoft Outlook Express versions 5.5 and 6.0. WebVPN should support other SMTPS, POP3S, or IMAP4S e-mail programs via port forwarding, such as Netscape Mail, Lotus Notes, and Eudora, but we have not verified them.
Using E-mail via Web Access	Web-based e-mail product installed	Supported: - Outlook Web Access For best results, use OWA on Internet Explorer 6.x or higher, Mozilla 1.7, or Firefox 1.x. - Lotus iNotes Other web-based e-mail products should also work, but we have not verified them.
Using E-mail via E-mail Proxy	SSL-enabled mail application installed Do not set the security appliance SSL version to TLSv1 Only. Outlook and Outlook Express do not support TLS.	Supported mail applications: - Microsoft Outlook - Microsoft Outlook Express versions 5.5 and 6.0 - Netscape Mail version 7 - Eudora 4.2 for Windows 2000 Other SSL-enabled mail clients should also work, but we have not verified them.
	Mail application configured	See instructions and examples for your mail application in “Using E-Mail over WebVPN.”

Capturing WebVPN Data

The CLI capture command lets you log information about websites that do not display properly over a WebVPN connection. This data can help your Cisco customer support engineer troubleshoot problems. The following sections describe how to use the capture command:

- [Creating a Capture File](#)
- [Using a Browser to Display Capture Data](#)



Note

Enabling WebVPN capture affects the performance of the security appliance. Be sure to disable the capture after you generate the capture files needed for troubleshooting.

Creating a Capture File

Perform the following steps to capture data about a WebVPN session to a file.

-
- Step 1** To start the WebVPN capture utility, use the **capture** command from privileged EXEC mode.
- ```
capture capture_name type webvpn user webvpn_username
```
- where:
- *capture\_name* is a name you assign to the capture, which is also prepended to the name of the capture files.
  - *webvpn\_user* is the username to match for capture.
- The capture utility starts.
- Step 2** A WebVPN user logs in to begin a WebVPN session. The capture utility is capturing packets. Stop the capture by using the **no** version of the command.
- ```
no capture capture_name
```
- The capture utility creates a *capture_name.zip* file, which is encrypted with the password **koleso**.
- Step 3** Send the .zip file to Cisco Systems, or attach it to a Cisco TAC service request.
- Step 4** To look at the contents of the .zip file, unzip it using the password **koleso**.
-

The following example creates a capture named *hr*, which captures WebVPN traffic for user2 to a file:

```
hostname# capture hr type webvpn user user2
WebVPN capture started.
  capture name    hr
  user name       user2
hostname# no capture hr
```

Using a Browser to Display Capture Data

Perform the following steps to capture data about a WebVPN session and view it in a browser.

-
- Step 1** To start the WebVPN capture utility, use the **capture** command from privileged EXEC mode.
- ```
capture capture_name type webvpn user webvpn_username
```
- where:
- *capture\_name* is a name you assign to the capture, which is also prepended to the name of the capture files.
  - *webvpn\_user* is the username to match for capture.
- The capture utility starts.
- Step 2** A WebVPN user logs in to begin a WebVPN session. The capture utility is capturing packets. Stop the capture by using the **no** version of the command.
- Step 3** Open a browser and in the address box enter
- ```
https://IP_address or hostname of the security appliance/webvpn_capture.html
```

The captured content displays in a sniffer format.

- Step 4** When you finish examining the capture content, stop the capture by using the **no** version of the command.
-



CHAPTER 38

Configuring SSL VPN Client

The SSL VPN Client (SVC) is a VPN tunneling technology that gives remote users the benefits of an IPSec VPN client without the need for network administrators to install and configure IPSec VPN clients on remote computers. The SVC uses the SSL encryption that is already present on the remote computer as well as the WebVPN login and authentication of the security appliance.

To establish an SVC session, the remote user enters the IP address of a WebVPN interface of the security appliance in the browser, and the browser connects to that interface and displays the WebVPN login screen. If the user satisfies the login and authentication, and the security appliance identifies the user as *requiring* the SVC, the security appliance downloads the SVC to the remote computer. If the security appliance identifies the user as having the *option* to use the SVC, the security appliance downloads the SVC to the remote computer while presenting a link on the user screen to skip the SVC installation.

After downloading, the SVC installs and configures itself, and then the SVC either remains or uninstalls itself (depending on the configuration) from the remote computer when the connection terminates.

This section covers the following topics:

- [Installing SVC, page 38-1](#)
- [Enabling SVC, page 38-3](#)
- [Enabling Permanent SVC Installation, page 38-4](#)
- [Enabling Rekey, page 38-5](#)
- [Enabling and Adjusting Dead Peer Detection, page 38-5](#)
- [Enabling Keepalive, page 38-6](#)
- [Using SVC Compression, page 38-6](#)
- [Viewing SVC Sessions, page 38-7](#)
- [Logging Off SVC Sessions, page 38-8](#)
- [Updating SVCs, page 38-8](#)

Installing SVC

This section presents the platform requirements and the procedure for installing SVC.

Platform Requirements

The SVC requires Windows 2000 or Windows XP on the remote computer.

Installing the SVC Software

Installing SVC consists of copying the SVC images to the security appliance and assigning an order to the images. Perform the following steps to install SVC:

- Step 1** Copy the SVC images to the security appliance using the **copy** command from privileged EXEC mode, or using another method. In this example, the images are copied from a tftp server using the **copy tftp** command:

```
hostname# copy tftp flash
Address or name of remote host []? 209.165.200.226
Source filename []? sslclient-win-1.0.2.127.pkg
Destination filename []? sslclient-win-1.0.2.127.pkg
Accessing tftp://209.165.200.226/sslclient-win-1.0.2.127.pkg...!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Writing file
disk0:/cdisk71...!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
319662 bytes copied in 3.695 secs (86511 bytes/sec)
```

- Step 2** Assign an order to the SVC images using the **svc image** command from webvpn mode:

svc image filename order

Numbering of the SVC images establishes the order in which the security appliance downloads them to the remote computer. It downloads the SVC image with the lowest number first. Therefore, you should assign the lowest number to the image used by the most commonly-encountered operating system.

In the following example, the output of the **show webvpn svc** command indicates that the windows.pkg image has an order number of 1, and the windows2.pkg image has an order number of 2. When a remote computer attempts to establish an SVC connection, the windows.pkg image downloads first. If the image does not match the operating system, the windows2.pkg image downloads:

```
hostname(config)# webvpn
hostname(config-webvpn)# show webvpn svc
1. disk0:/windows.pkg 1
   CISCO STC win2k+ 1.0.0
   1,0,2,132
   Thu 08/25/2005 21:51:30.43

2. disk0:/windows2.pkg 2
   CISCO STC win2k+ 1.0.0
   1,0,0,164
   Thu 02/17/2005 20:09:22.43

2 SSL VPN Client(s) installed
```

Then the SVC archive images are re-ordered using the **svc image** command, with the windows2.pkg image as the first image downloaded to the remote PC, and the windows.pkg image downloaded second:

```
hostname(config-webvpn)# no svc image
hostname(config-webvpn)# svc image windows2.pkg 1
hostname(config-webvpn)# svc image windows.pkg 2
```

Reentering the **show webvpn svc** command shows the new order of the images:

```
hostname(config-webvpn)# show webvpn svc
1. disk0:/windows2.pkg 1
   CISCO STC win2k+ 1.0.0
   1,0,2,132
   Thu 08/25/2005 21:51:30.43

2. disk0:/windows.pkg 2
   CISCO STC win2k+ 1.0.0
```

```
1,0,0,164
Thu 02/17/2005 20:09:22.43

2 SSL VPN Client(s) installed
```

Enabling SVC

After installing SVC, you can enable SVC by performing the following steps:

-
- Step 1** Enable WebVPN on an interface using the **enable** command from webvpn mode:
- enable interface**
- For example:
- ```
hostname(config)# webvpn
hostname(config-webvpn)# enable outside
```
- Step 2** From webvpn mode, enter the **svc enable** command to enable the security appliance to download SVC images to remote computers:
- svc enable**
- For example:
- ```
hostname(config-webvpn)# svc enable
```
- Step 3** Configure a method of address assignment. You can use DHCP, and/or user-assigned addressing. You can also create a local IP address pool using the **ip local pool** command from webvpn mode:
- ip local pool poolname startaddr-endaddr mask mask**
- The following example creates the local IP address pool *vpn_users*:
- ```
hostname(config-webvpn)# ip local pool vpn_users 209.165.200.225-209.165.200.254
mask 255.255.255.224
```
- Step 4** Assign IP addresses to a tunnel group. One method you can use to do this is to configure a local IP address pool with the **address-pool** command from general-attributes mode:
- address-pool poolname**
- To do this, first enter the **tunnel-group name general-attributes** command to enter general-attributes mode. Then specify the local IP address pool using the **address-pool** command.
- In the following example, the user configures the existing tunnel group *telecommuters* to use the address pool *vpn\_users* created in step 3:
- ```
hostname(config)# tunnel-group telecommuters general-attributes
hostname(config-tunnel-general)# address-pool vpn_users
```
- Step 5** Assign a default group policy to the tunnel group with the **default-group-policy** command from tunnel group general attributes mode:
- default-group-policy name**
- In the following example, the user assigns the group policy *sales* to the tunnel group *telecommuters*:
- ```
hostname(config-tunnel-general)# default-group-policy sales
```
- Step 6** Create and enable a group alias that displays in the group list on the WebVPN Login page using the **group-alias** command from tunnel group webvpn attributes mode:
- group-alias name enable**

First exit to global configuration mode, and then enter the **tunnel-group name webvpn-attributes** command to enter tunnel group webvpn attributes mode.

In the following example, the user enters webvpn attributes configuration mode for the tunnel group *telecommuters*, and creates the group alias *sales\_department*:

```
hostname(config)# tunnel-group telecommuters webvpn-attributes
hostname(config-tunnel-webvpn)# group-alias sales_department enable
```

- Step 7** Enable the display of the tunnel-group list on the WebVPN Login page from webvpn mode:

**tunnel-group-list enable**

First exit to global configuration mode, and then enter webvpn mode.

In the following example, the user enters webvpn mode, and then enables the tunnel group list:

```
hostname(config)# webvpn
hostname(config-webvpn)# tunnel-group-list enable
```

- Step 8** Identify WebVPN as a permitted VPN tunneling protocol for the group or user with the **vpn-tunnel-protocol webvpn** command in group-policy mode or username mode:

**vpn-tunnel-protocol webvpn**

To do this, first exit to global configuration mode, enter the **group-policy name attributes** command to enter group-policy mode, or the **username name attributes** command to enter username mode, and then enter the **webvpn** command to enter webvpn mode and change the WebVPN settings for the group or user.

The following example identifies WebVPN as a permitted tunneling protocol for the group-policy *sales*:

```
hostname(config)# group-policy sales attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# vpn-tunnel-protocol webvpn
```

- Step 9** Enable or require an SVC for a specific group or user by using the **svc** command from either group-policy webvpn mode or username webvpn mode:

**svc {none | enable | required}**

The following example sets the SVC to *required* for the existing group-policy *sales*:

```
hostname(config-group-webvpn)# svc required
```

For more information about assigning users to group policies, see [Chapter 30, “Configuring Tunnel Groups, Group Policies, and Users”](#).

## Enabling Permanent SVC Installation

Enabling permanent SVC installation disables the automatic uninstalling feature of the SVC. The SVC remains installed on the remote computer for subsequent SVC connections, reducing the SVC connection time for the remote user.

To enable permanent SVC installation for a specific group or user, use the **svc keep-installer** command from group-policy or username webvpn modes:

**svc keep-installer {installed | none}**  
**no svc keep-installer {installed | none}**

**Where:**

**installed** specifies the SVC is permanently installed on the remote computer.

**none** specifies the SVC is removed from the remote computer after the active SVC connection terminates.

The default is that permanent installation of the SVC is disabled. The SVC on the remote computer uninstalls at the end of every SVC session.

The following example configures the existing group-policy *sales* to keep the SVC installed on the remote computer:

```
hostname(config)# group-policy sales attributes
hostname(config-group-policy)# webvpn
hostname(config-group-policy)# svc keep-installer installed
```

## Enabling Rekey

When the security appliance and the SVC perform a rekey, they renegotiate the crypto keys and initialization vectors, increasing the security of the connection.

To enable the SVC to perform a rekey on an SVC session for a specific group or user, use the **svc rekey** command from group-policy and username webvpn modes.

```
svc rekey {method {new-tunnel | none | ssl} | time minutes}
no svc rekey {method {new-tunnel | none | ssl} | time minutes}
```

Where:

**method new-tunnel** specifies that the SVC establishes a new tunnel during SVC rekey.

**method none** disables SVC rekey.

**method ssl** specifies that SSL renegotiation takes place during SVC rekey.

**time minutes** specifies the number of minutes from the start of the session until the rekey takes place, from 1 to 10080 (1 week).

In the following example, the SVC is configured to renegotiate with SSL during rekey, which takes place 30 minutes after the session begins, for the existing group-policy *sales*:

```
hostname(config)# group-policy sales attributes
hostname(config-group-policy)# webvpn
hostname(config-group-policy)# svc rekey method ssl
hostname(config-group-policy)# svc rekey time 30
```

## Enabling and Adjusting Dead Peer Detection

Dead Peer Detection (DPD) ensures that the security appliance (gateway) or the SVC can quickly detect a condition where the peer is not responding, and the connection has failed.

To enable DPD on the security appliance or SVC for a specific group or user, and to set the frequency with which either the security appliance or SVC performs DPD, use the **svc dpd-interval** command from group-policy or username webvpn mode:

```
svc dpd-interval {[gateway {seconds | none}] | [client {seconds | none}]}
no svc dpd-interval {[gateway {seconds | none}] | [client {seconds | none}]}
```

Where:

**gateway seconds** enables DPD performed by the security appliance (gateway) and specifies the frequency, from 30 to 3600 seconds, with which the security appliance (gateway) performs DPD.

**gateway none** disables DPD performed by the security appliance.

**client seconds** enable DPD performed by the SVC (client), and specifies the frequency, from 30 to 3600 seconds, with which the SVC performs DPD.

**client none** disables DPD performed by the SVC.

To remove the **svc dpd-interval** command from the configuration, use the **no** form of the command:

The following example sets the frequency of DPD performed by the security appliance to 3000 seconds, and the frequency of DPD performed by the SVC set to 1000 seconds for the existing group-policy *sales*:

```
hostname(config)# group-policy sales attributes
hostname(config-group-policy)# webvpn
hostname(config-group-policy)# svc dpd-interval gateway 3000
hostname(config-group-policy)# svc dpd-interval client 1000
```

## Enabling Keepalive

You can adjust the frequency of keepalive messages to ensure that an SVC connection through a proxy, firewall, or NAT device remains open, even if the device limits the time that the connection can be idle. Adjusting the frequency also ensures that the SVC does not disconnect and reconnect when the remote user is not actively running a socket-based application, such as Microsoft Outlook or Microsoft Internet Explorer.

To set the frequency of keepalive messages, use the **svc keepalive** command from group-policy or username webvpn modes:

```
svc keepalive {none | seconds}
no svc keepalive {none | seconds}
```

**Where:**

**none** disables SVC keepalive messages.

**seconds** enables the SVC to send keepalive messages, and specifies the frequency of the messages in the range of 15 to 600 seconds.

The default is keepalive messages are disabled.

Use the **no** form of the command to remove the command from the configuration and cause the value to be inherited:

In the following example, the security appliance is configured to enable the SVC to send keepalive messages with a frequency of 300 seconds (5 minutes), for the existing group-policy *sales*:

```
hostname(config)# group-policy sales attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# svc keepalive 300
```

## Using SVC Compression

SVC compression increases the communications performance between the security appliance and the SVC by reducing the size of the packets being transferred. By default, compression for all SVC connections is enabled on the security appliance, both at the global level and for specific groups or users.

SVC compression can be set globally using the **compression svc** command from global configuration mode. It can also be set for specific groups or users with the **svc compression** command in group-policy and username webvpn modes. The global setting overrides the group-policy and username settings.

### Changing SVC Compression Globally

To change the global SVC compression settings, use the **compression svc** command from global configuration mode:

```
compression svc
```

```
no compression svc
```

To remove the command from the configuration, use the **no** form of the command.

In the following example, compression is disabled for all SVC connections globally:

```
hostname(config)# no compression svc
```

### Changing SVC Compression for Groups and Users

To change compression for a specific group or user, use the **svc compression** command in the group-policy and username webvpn modes:

```
svc compression {deflate | none}
```

```
no svc compression {deflate | none}
```

By default, for groups and users, SVC compression is set to *deflate* (enabled).

To remove the **svc compression** command from the configuration and cause the value to be inherited from the global setting, use the **no** form of the command:

In the following example, SVC compression is disabled for the group-policy sales:

```
hostname(config)# group-policy sales attributes
hostname(config-group-policy)# webvpn
hostname(config-group-webvpn)# svc compression none
```



#### Note

The **compression svc** command configured from global configuration mode overrides the **svc compression** command configured in group-policy and username webvpn modes.

## Viewing SVC Sessions

You can view information about active SVC sessions using the **show vpn-sessiondb** command in privileged EXEC mode:

```
show vpn-sessiondb svc
```

The following example shows the output of the **show vpn-sessiondb svc** command:

```
hostname# show vpn-sessiondb svc

Session Type: SSL VPN Client

Username : lee
Index : 1
Protocol : SSL VPN Client
Hashing : SHA1
TCP Dst Port : 443
Bytes Tx : 20178
Pkts Tx : 27
Client Ver : Cisco STC 1.1.0.117
Client Type : Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0; Q312461)
Group : DfltGrpPolicy
Login Time : 14:32:03 UTC Wed Apr 20 2005
Duration : 0h:00m:04s

IP Addr : 161.44.128.249
Encryption : 3DES
Auth Mode : userPassword
TCP Src Port : 54230
Bytes Rx : 8662
Pkts Rx : 19
```

Filter Name :

## Logging Off SVC Sessions

To log off all SVC sessions, use the **vpn-sessiondb logoff svc** command in global configuration mode:

```
vpn-sessiondb logoff svc
```

The following example logs off all SVC sessions:

```
hostname# vpn-sessiondb logoff svc
INFO: Number of sessions of type "svc" logged off : 1
```

You can log off individual SVC sessions using either the **name option**, or the **index option**:

```
vpn-session-db logoff name name
```

```
vpn-session-db logoff index index
```

You can find both the username and the index number (established by the order of the SVC images) in the output of the **show vpn-sessiondb svc** command. The following example shows the username *lee* and index number *1*.

```
hostname# show vpn-sessiondb svc

Session Type: SSL VPN Client

Username : lee
Index : 1
Protocol : SSL VPN Client
Hashing : SHA1
TCP Dst Port : 443
Bytes Tx : 20178
Pkts Tx : 27
Client Ver : Cisco STC 1.1.0.117
Client Type : Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0; Q312461)
Group : DfltGrpPolicy
Login Time : 14:32:03 UTC Wed Apr 20 2005
Duration : 0h:00m:04s
Filter Name :
```

The following example terminates the session using the **name** option of the **vpn-session-db logoff** command:

```
hostname# vpn-session-db logoff name lee
INFO: Number of sessions with name "lee" logged off : 1
```

## Updating SVCs

You can update the SVC images on the security appliance at any time using the following procedure:

- 
- Step 1** Copy the new SVC images to the security appliance using the **copy** command from privileged EXEC mode, or using another method.
  - Step 2** If the new SVC image files have the same filenames as the files already loaded, reenter the **svc image** command that is in the configuration. If the new filenames are different, uninstall the old files using the **no svc image** command. Then use the **svc image** command to assign an order to the SVC images and cause the security appliance to load the new SVC images.





## CHAPTER 39

# Configuring Certificates

---

This chapter describes how to configure certificates. CAs are responsible for managing certificate requests and issuing digital certificates. A digital certificate contains information that identifies a user or device. Some of this information can include a name, serial number, company, department, or IP address. A digital certificate also contains 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.

This chapter includes the following sections:

- [Public Key Cryptography, page 39-1](#)
- [Certificate Configuration, page 39-5](#)

## Public Key Cryptography

This section includes the following topics:

- [About Public Key Cryptography, page 39-1](#)
- [Certificate Scalability, page 39-2](#)
- [About Key Pairs, page 39-2](#)
- [About Trustpoints, page 39-3](#)
- [About CRLs, page 39-3](#)
- [Supported CA Servers, page 39-5](#)

## About Public Key Cryptography

Digital signatures, enabled by public key cryptography, provide a means 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 having 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 out-of-band or through an operation done at installation. For instance, 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 every peer with which it communicates, and every new peer you add to a network would thus require a configuration change on every peer with which you need it to communicate securely.

When you use digital certificates, each peer is enrolled with a CA. When two peers attempt 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 and each certificate is authenticated by the CA, and all participating peers recognize the CA as an authenticating authority. This 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.

## About Key Pairs

Key pairs are RSA keys.

- 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. The default size is 1024.
- For signature operations, the supported maximum key size is 4096 bits.
- You can generate a general purpose RSA key pair, used for both signing and encryption, or you can generate separate RSA key pairs for each purpose.

Separate signing and encryption keys helps reduce exposure of the keys. This is because SSL uses a key for encryption but not signing but IKE uses a key for signing but not encryption. By using separate keys for each, exposure of the keys is minimized.

## About Trustpoints

Trustpoints let you manage and track CAs and certificates. A trustpoint is a representation of a CA or identity pair. A trustpoint contains 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 a security appliance has multiple trustpoints that share the same CA, only one of these trustpoints sharing the CA can be used to validate user certificates. Use the **support-user-cert-validation** command to control which trustpoint sharing a CA is used for validation of user certificates issued by that CA.

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 is useful if you wish to manually duplicate a trustpoint configuration on a different security appliance.

## About 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, due to security concerns or a change of name or association. CAs periodically issue a signed list of revoked certificates. Enabling revocation checking forces the security appliance to check that the CA has not revoked a certificate every time it uses that certificate for authentication.

When you enable revocation checking, during the PKI certificate validation process the security appliance checks certificate revocation status. It can use either CRL checking or Online Certificate Status Protocol or both, with the second method you set in effect only when the first method returns an error, for example, that the server is unavailable.

With CRL checking, the security appliance retrieves, parses, and caches Certificate Revocation Lists, which provide a complete list of revoked certificates. OCSP offers a more scalable method of checking revocation status in that it localizes certificate status on a Validation Authority, which it queries for the status of a specific certificate.

## About CRLs

Certificate Revocation Lists provide the security appliance with one means of determining whether a certificate that is within its valid time range has been revoked by its issuing CA. CRL configuration is a part of the configuration of a trustpoint.

You can configure the security appliance to make CRL checks mandatory when authenticating a certificate (**revocation-check crl** command). You can also make the CRL check optional by adding the **none** argument (**revocation-check crl none** command), which allows the certificate authentication to succeed when the CA is unavailable to provide updated CRL data.

The security appliance can retrieve CRLs from CAs using HTTP, SCEP, or LDAP. CRLs retrieved for each trustpoint are cached for a length of time configurable for each trustpoint.

When the security appliance has cached a CRL for more than the length of time it is configured to cache CRLs, the security appliance considers the CRL too old to be reliable, or “stale”. The security appliance attempts to retrieve a newer version of the CRL the next time a certificate authentication requires checking the stale CRL.

The security appliance caches CRLs for a length 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 security appliance requires and uses the NextUpdate field with the **enforcenextupdate** command.

The security appliance uses these two factors as follows:

- If the NextUpdate field is not required, the security appliance marks CRLs as stale after the length of time defined by the **cache-time** command.
- If the NextUpdate field is required, the security appliance 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 security appliance marks CRLs as stale in 70 minutes.

If the security appliance 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.

For information about configuring CRL behavior for a trustpoint, see the [“Configuring CRLs for a Trustpoint” section on page 39-13](#).

## About OCSP

Online Certificate Status Protocol provides the security appliance with a means of determining whether a certificate that is within its valid time range has been revoked by its issuing CA. OCSP configuration is a part of the configuration of a trustpoint.

OCSP localizes certificate status on a Validation Authority (an OCSP server, also called the *responder*) which the security appliance queries for the status of a specific certificate. It provides better scalability and more up-to-date revocation status than does CRL checking. It helps organizations with large PKI installations deploy and expand secure networks.

You can configure the security appliance to make OCSP checks mandatory when authenticating a certificate (**revocation-check ocs**p command). You can also make the OCSP check optional by adding the **none** argument (**revocation-check ocs**p **none** command), which allows the certificate authentication to succeed when the Validation Authority is unavailable to provide updated OCSP data.

Our implementation of OCSP provides three ways to define the OCSP server URL. The security appliance uses these servers in the following order:

1. The OCSP URL defined in a match certificate override rule (**match certificate** command).
2. The OCSP URL configured in the **ocsp url** command.
3. The AIA field of the client certificate.



### Note

To configure a trustpoint to validate a self-signed OCSP responder certificate, you import the self-signed responder certificate into its own trustpoint as a trusted CA certificate. Then you configure the **match certificate** command in the client certificate validating trustpoint to use the trustpoint that contains the self-signed OCSP responder certificate to validate the responder certificate. The same applies for configuring validating responder certificates external to the validation path of the client certificate.

The OCSP server (responder) certificate typically signs the OCSP response. After receiving the response, the security appliance tries to verify the responder certificate. The CA normally sets the lifetime of its OCSP responder certificate to a relatively short period to minimize the chance of it being compromised. The CA typically also includes an `ocsp-no-check` extension in the responder certificate indicating that this certificate does not need revocation status checking. But if this extension is not present, the security appliance tries to check its revocation status using the same method specified in the trustpoint. If the responder certificate is not verifiable, revocation checks fails. To avoid this possibility, configure **revocation-check none** in the responder certificate validating trustpoint, while configuring **revocation-check ocsp** for the client certificate.

---

## Supported CA Servers

The security appliance supports the following CA servers:

- Cisco IOS CS
- Baltimore Technologies
- Entrust
- Microsoft Certificate Services
- Netscape CMS
- RSA Keon
- VeriSign

## Certificate Configuration

This section describes how to configure the security appliance with certificates and other procedures related to certificate use and management.

This section includes the following topics:

- [Preparing for Certificates, page 39-5](#)
- [Configuring Key Pairs, page 39-6](#)
- [Configuring Trustpoints, page 39-7](#)
- [Obtaining Certificates, page 39-9](#)
- [Configuring CRLs for a Trustpoint, page 39-13](#)
- [Exporting and Importing Trustpoints, page 39-14](#)
- [Configuring CA Certificate Map Rules, page 39-15](#)

## Preparing for Certificates

Before you configure a security appliance with certificates, ensure that the security appliance is configured properly to support certificates. An improperly configured security appliance can cause enrollment to fail or for enrollment to request a certificate containing inaccurate information.

To prepare a security appliance for certificates, perform the following steps:

- 
- Step 1** Ensure that the hostname and domain name of the security appliance are configured correctly. You can use the **show running-config** command to view the hostname and domain name as currently configured. For information about configuring the hostname, see the [“Setting the Hostname” section on page 8-2](#). For information about configuring the domain name, see the [“Setting the Domain Name” section on page 8-2](#).
- Step 2** Be sure that the security appliance clock is set accurately before configuring the CA. Certificates have a date and time that they become valid and that they expire. When the security appliance enrolls with a CA and gets a certificate, the security appliance checks that the current time is within the valid range for the certificate. If it is outside that range, enrollment fails. For information about setting the clock, see the [“Setting the Date and Time” section on page 8-2](#).
- 

## Configuring Key Pairs

This section includes the following topics:

- [Generating Key Pairs, page 39-6](#)
- [Removing Key Pairs, page 39-7](#)

## Generating Key Pairs

Key pairs are RSA keys, as discussed in the [“About Key Pairs” section on page 39-2](#). You must generate key pairs for the types of certification you want to use.

To generate key pairs, perform the following steps:

- 
- Step 1** Generate the types of key pairs needed for your PKI implementation. To do so, perform the following steps, as applicable:
- a. If you want to generate RSA key pairs, use the **crypto key generate rsa** command.  

```
hostname/contexta(config)# crypto key generate rsa
```

If you do not use additional keywords this command generates one general purpose RSA key pair. Because the key modulus is not specified, the default key modulus of 1024 is used. You can specify other modulus sizes with the **modulus** keyword. You can also assign a label to each key pair using the **label** keyword. The label is referenced by the trustpoint that uses the key pair. If you do not assign a label, the key pair is automatically labeled <Default-RSA-Key>.

```
hostname/contexta(config)# crypto key generate rsa label key-pair-label
```
- Step 2** (Optional) Use the **show crypto key mypubkey** command to view key pair(s). The following example shows an RSA general-purpose key:
- ```
hostname/contexta(config)# show crypto key mypubkey  
Key pair was generated at: 16:39:47 central Feb 10 2005  
Key name: <Default-RSA-Key>  
Usage: General Purpose Key  
Modulus Size (bits): 1024  
Key Data:
```

```

30819f30 0d06092a 864886f7 0d010101 05000381 8d003081 89028181 00ea51b7
0781848f 78bccac2 4a1b5b8d 2f3e30b4 4cae9f86 f4485207 159108c9 f5e49103
9eeb0f5d 45fd1811 3b4aafce 292b3b64 b4124a6f 7a777b08 75b88df1 8092a9f8
5508e9e5 2c271245 7fd1c0c3 3aaf1e04 c7c4efa4 600f4c4a 6afe56ad c1d2c01c
e08407dd 45d9e36e 8cc0bfef 14f9e6ac eca141e4 276d7358 f7f50d13 79020301 0001
Key pair was generated at: 16:34:54 central Feb 10 2005

```

- Step 3** Save the key pair you have generated. To do so, save the running configuration by entering the **write memory** command.

Removing Key Pairs

To remove key pairs, use the **crypto key zeroize** command in global configuration mode.

The following example removes RSA key pairs:

```

hostname(config)# crypto key zeroize rsa
WARNING: All RSA keys will be removed.
WARNING: All device certs issued using these keys will also be removed.

Do you really want to remove these keys? [yes/no] y
hostname(config)#

```

Configuring Trustpoints

For information about trustpoints, see the [“About Trustpoints” section on page 39-3](#).

To configure a trustpoint, perform the following steps:

- Step 1** Create a trustpoint corresponding to the CA from which the security appliance needs to receive its certificate.

```
hostname/contexta(config)# crypto ca trustpoint trustpoint
```

For example, to declare a trustpoint called Main:

```
hostname/contexta(config)# crypto ca trustpoint Main
hostname/contexta(config-ca-trustpoint)#

```

Upon entering this command, you enter the Crypto ca trustpoint configuration mode.

- Step 2** Specify the enrollment method to be used with this trustpoint.

To specify the enrollment method, do one of the following items:

- To specify SCEP enrollment, use the **enrollment url** command to configure the URL to be used for SCEP enrollment with the trustpoint you declared. For example, if the security appliance requests certificates from trustpoint Main using the URL `http://10.29.67.142:80/certsrv/mscep/mscep.dll`, then the command would be as follows:

```
hostname/contexta(config-ca-trustpoint)# enrollment url
http://10.29.67.142:80/certsrv/mscep/mscep.dll

```

- To specify manual enrollment, use the **enrollment terminal** command to indicate that you will paste the certificate received from the CA into the terminal.

Step 3 As needed, specify other characteristics for the trustpoint. The characteristics you need to define depend upon your CA and its configuration. You can specify characteristics for the trustpoint using the following commands. Refer to the *Cisco Security Appliance Command Reference* for complete descriptions and usage guidelines of these commands.

- **accept-subordinates**—Indicates whether CA certificates subordinate to the CA associated with the trustpoint are accepted if delivered during phase one IKE exchange when not previously installed on the device.
- **crl required | optional | nocheck**—Specifies CRL configuration options. When you enter the **crl** command with the **optional** keyword included within the command statement, certificates from peers can still be accepted by your security appliance even if the CRL is not accessible to your security appliance.



Note If you chose to enable required or optional CRL checking, be sure you configure the trustpoint for CRL management, which should be completed after you have obtained certificates. For details about configuring CRL management for a trustpoint, see the [“Configuring CRLs for a Trustpoint”](#) section on page 39-13.

- **crl configure**—Enters CRL configuration mode.
- **default enrollment**—Returns all enrollment parameters to their system default values. Invocations of this command do not become part of the active configuration.
- **email address**—During enrollment, asks the CA to include the specified email address in the Subject Alternative Name extension of the certificate.
- **enrollment retry period**—(Optional) Specifies a retry period in minutes. This characteristic only applies if you are using SCEP enrollment.
- **enrollment retry count**—(Optional) Specifies a maximum number of permitted retries. This characteristic only applies if you are using SCEP enrollment.
- **enrollment terminal**—Specifies cut and paste enrollment with this trustpoint.
- **enrollment url URL**—Specifies automatic enrollment (SCEP) to enroll with this trustpoint and configures the enrollment URL.
- **fqdn fqdn**—During enrollment, asks the CA to include the specified fully qualified domain name in the Subject Alternative Name extension of the certificate.
- **id-cert-issuer**—Indicates whether the system accepts peer certificates issued by the CA associated with this trustpoint.
- **ip-address ip-address**—During enrollment, asks the CA to include the IP address of the security appliance in the certificate.
- **keypair name**—Specifies the key pair whose public key is to be certified.
- **match certificate map**—Configures OCSP URL overrides and trustpoints to use to validate OCSP responder certificates
- **ocsp disable-nonce**—Disable the nonce extension on an OCSP request; the nonce extension cryptographically binds requests with responses to avoid replay attacks.
- **ocsp url**—Configures an OCSP server for the security appliance to use to check all certificates associated with a trustpoint rather than the server specified in the AIA extension of the client certificate.
- **password string**—Specifies a challenge phrase that is registered with the CA during enrollment. The CA typically uses this phrase to authenticate a subsequent revocation request.

- **revocation-check**—Sets one or more methods for revocation checking: CRL, OCSP, and none.
- **subject-name** *X.500 name*—During enrollment, asks the CA to include the specified subject DN in the certificate.
- **serial-number**—During enrollment, asks the CA to include the security appliance serial number in the certificate.
- **support-user-cert-validation**—If enabled, the configuration settings to validate a remote user certificate can be taken from this trustpoint, provided that this trustpoint is authenticated to the CA that issued the remote certificate.
- **exit**—Leaves the mode.

Step 4 Save the trustpoint configuration. To do so, save the running configuration by entering the **write memory** command.

Obtaining Certificates

The security appliance 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 security appliance needs two certificates, one for each purpose. In other key configurations, only one certificate is needed.

The security appliance supports 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 security appliance. For remote access VPNs, you must enroll each security appliance and each remote access VPN client.

This section includes the following topics:

- [Obtaining Certificates with SCEP, page 39-9](#)
- [Obtaining Certificates Manually, page 39-11](#)

Obtaining Certificates with SCEP

This procedure provides steps for configuring certificates using SCEP. Repeat these steps for each trustpoint you configure for automatic enrollment. When you have completed this procedure, the security appliance will have received a CA certificate for the trustpoint and one or two certificates for signing and encryption purposes. If you use general-purpose RSA keys, the certificate received is for signing and encryption. If you use separate RSA keys for signing and encryption, the security appliance receives separate certificates for each purpose.



Note

Whether a trustpoint uses SCEP for obtaining certificates is determined by the use of the **enrollment url** command when you configure the trustpoint (see the [“Configuring Trustpoints” section on page 39-7](#)).

To obtain certificates with SCEP, perform the following steps:

Step 1 Obtain the CA certificate for the trustpoint you configured.

```
hostname/contexta(config)# crypto ca authenticate trustpoint
```

For example, using trustpoint named Main, which represents a subordinate CA:

```
hostname/contexta(config)# crypto ca authenticate Main
```

```
INFO: Certificate has the following attributes:
Fingerprint:      3736ffc2 243ecf05 0c40f2fa 26820675
Do you accept this certificate? [yes/no]: y
```

```
Trustpoint 'Main' is a subordinate CA and holds a non self signed cert.
Trustpoint CA certificate accepted.
```

Step 2 Enroll the security appliance with the trustpoint. This process retrieves a certificate for signing data and, depending upon the type of keys you configured, for encrypting data.

Step 3 To perform enrollment, use the **crypto ca enroll** command. Before entering this command, contact your CA administrator because the administrator may need to authenticate your enrollment request manually before the CA grants its certificates.

```
hostname(config)# crypto ca enroll trustpoint
```

If the security appliance does not receive a certificate from the CA within 1 minute (the default) of sending a certificate request, it resends the certificate request. The security appliance continues sending a certificate request every 1 minute until a certificate is received.



Note If the fully qualified domain name configured for the trustpoint is not identical to the fully qualified domain name of the security appliance, including the case of the characters, a warning appears. If needed, you can exit the enrollment process, make any necessary corrections, and enter the **crypto ca enroll** command again.

The following enrollment example performs enrollment with the trustpoint named Main:

```
hostname(config)# crypto ca enroll Main
%
% Start certificate enrollment ..
% Create a challenge password. You will need to verbally provide this
% password to the CA Administrator in order to revoke your certificate.
% For security reasons your password will not be saved in the configuration.
% Please make a note of it.
Password: 2b0rn0t2b
Re-enter password: 2b0rn0t2b
% The subject name in the certificate will be: securityappliance.example.com
% The fully-qualified domain name in the certificate will be:
securityappliance.example.com
% Include the device serial number in the subject name? [yes/no]: no
Request certificate from CA [yes/no]: yes
% Certificate request sent to Certificate authority.
```



Note The password is required if the certificate for the security appliance needs to be revoked, so it is crucial that you remember this password. Note it and store it in a safe place.

You must enter the **crypto ca enroll** command for each trustpoint with which the security appliance needs to enroll.



Note If your security appliance reboots after you issued the **crypto ca enroll** command but before you received the certificate, reissue the **crypto ca enroll** command and notify the CA administrator.

- Step 4** Verify that the enrollment process was successful using the **show crypto ca certificate** command. For example, to show the certificate received from trustpoint Main:

```
hostname/contexta(config)# show crypto ca certificate Main
```

The output of this command shows the details of the certificate issued for the security appliance and the CA certificate for the trustpoint.

- Step 5** Save the configuration using the **write memory** command:

```
hostname/contexta(config)# write memory
```

Obtaining Certificates Manually

This procedure provides steps for configuring certificates using manual certificate requests. Repeat these steps for each trustpoint you configure for manual enrollment. When you have completed this procedure, the security appliance will have received a CA certificate for the trustpoint and one or two certificates for signing and encryption purposes. If you use general-purpose RSA keys, the certificate received is for signing and encryption. If you use separate RSA keys for signing and encryption, the certificates received are used for each purpose exclusively.



Note

Whether a trustpoint requires that you manually obtain certificates is determined by the use of the **enrollment terminal** command when you configure the trustpoint (see the [“Configuring Trustpoints” section on page 39-7](#)).

To obtain certificates manually, perform the following steps:

- Step 1** Obtain a base-64 encoded CA certificate from the CA represented by the trustpoint.
- Step 2** Import the CA certificate. To do so, use the **crypto ca authenticate** command. The following example shows a CA certificate request for the trustpoint Main.

```
hostname (config)# crypto ca authenticate Main
Enter the base 64 encoded CA certificate.
End with a blank line or the word "quit" on a line by itself
MIIDRTCCAu+gAwIBAgIQKVcQP/KW74VP0NZzL+JbRTANBgqhkiG9w0BAQUFADCB
[ certificate data omitted ]
/7QEM8izy0EOTSErKu7Nd76jwf5e4qttkQ==
quit
```

```
INFO: Certificate has the following attributes:
Fingerprint:      24b81433 409b3fd5 e5431699 8d490d34
Do you accept this certificate? [yes/no]: y
Trustpoint CA certificate accepted.
```

```
% Certificate successfully imported
hostname (config)#
```

- Step 3** Generate a certificate request. To do so, use the **crypto ca enroll** command. The following example shows a certificate and encryption key request for the trustpoint Main, which is configured to use manual enrollment and general-purpose RSA keys for signing and encryption.

```
hostname (config)# crypto ca enroll Main
% Start certificate enrollment ..
```

```
% The fully-qualified domain name in the certificate will be:
securityappliance.example.com

% Include the device serial number in the subject name? [yes/no]: n

Display Certificate Request to terminal? [yes/no]: y
Certificate Request follows:

MIIBOCCAQkCAQAwIzEhMB8GCSqGSIb3DQEJAhYSRmVyYWxQaXguY2l2Y28uY29t
[ certificate request data omitted ]
jF4waw68eOxQxVmdgMWeQ+RbIOYmvt8g6hnBTrd0GdqjjVLt

---End - This line not part of the certificate request---

Redisplay enrollment request? [yes/no]: n
hostname (config)#
```



Note If you use separate RSA keys for signing and encryption, the **crypto ca enroll** command displays two certificate requests, one for each key. To complete enrollment, acquire a certificate for all certificate requests generated by the **crypto ca enroll** command.

Step 4 For each request generated by the **crypto ca enroll** command, obtain a certificate from the CA represented by the applicable trustpoint. Be sure the certificate is in base-64 format.

Step 5 For each certificate you receive from the CA, use the **crypto ca import certificate** command. The security appliance prompts you to paste the certificate to the terminal in base-64 format.



Note If you use separate RSA key pairs for signing and encryption, perform this step for each certificate separately. The security appliance determines automatically whether the certificate is for the signing or encryption key pair. The order in which you import the two certificates is irrelevant.

The following example manually imports a certificate for the trustpoint Main:

```
hostname (config)# crypto ca import Main certificate
% The fully-qualified domain name in the certificate will be:
securityappliance.example.com

Enter the base 64 encoded certificate.
End with a blank line or the word "quit" on a line by itself
[ certificate data omitted ]
quit
INFO: Certificate successfully imported
hostname (config)#
```

Step 6 Verify that the enrollment process was successful using the **show crypto ca certificate** command. For example, to show the certificate received from trustpoint Main:

```
hostname/contexta(config)# show crypto ca certificate Main
```

The output of this command shows the details of the certificate issued for the security appliance and the CA certificate for the trustpoint.

Step 7 Save the configuration using the **write memory** command:

```
hostname/contexta(config)# write memory
```

Configuring CRLs for a Trustpoint

If you want to use mandatory or optional CRL checking during certificate authentication, you must perform CRL configuration for each trustpoint. For more information about CRLs, see the [“About CRLs” section on page 39-3](#).

To configure CRLs for a trustpoint, perform the following steps:

-
- Step 1** Enter Crypto ca trustpoint configuration mode for the trustpoint whose CRL configuration you want to modify. To do so, enter the **crypto ca trustpoint** command.
- Step 2** If you have not already enabled CRLs, you can do so now by using the **crl** command with either the **required** or **optional** keyword. If you specify the **required** keyword, certificate authentication with this trustpoint cannot succeed if the CRL is unavailable.
- Step 3** Enter the **crl configure** command.

```
hostname/contexta(config-ca-trustpoint)# crl configure
hostname/contexta(config-ca-crl)#
```

Upon entering this command, you enter the crl configuration mode for the current trustpoint.



Tip To set all CRL configuration options to their default values, use the **default** command. At any time while performing CRL configuration, if you want to start over, enter this command and restart this procedure.

- Step 4** Configure the retrieval policy with the **policy** command. The following keywords for this command determine the policy.

- **cdp**—CRLs are retrieved only from the CRL distribution points specified in authenticated certificates.



Note SCEP retrieval is not supported by distribution points specified in certificates.

- **static**—CRLs are retrieved only from URLs you configure.
- **both**—CRLs are retrieved from CRL distribution points specified in authenticated certificates and from URLs you configure.

- Step 5** If you used the keywords static or both when you configured the CRL policy, you need to configure URLs for CRL retrieval, using the **url** command. You can enter up to 5 URLs, ranked 1 through 5.

```
hostname/contexta(config-ca-crl)# url n URL
```

where *n* is the rank assigned to the URL. To remove a URL, use the **no url *n*** command.

- Step 6** Configure the retrieval method with the **protocol** command. The following keywords for this command determine the retrieval method.

- **http**—Specifies HTTP as the CRL retrieval method.
- **ldap**—Specifies LDAP as the CRL retrieval method.
- **scep**—Specifies SCEP as the CRL retrieval method.

- Step 7** Configure how long the security appliance caches CRLs for the current trustpoint. To specify the number of minutes the security appliance waits before considering a CRL stale, enter the following command.

```
hostname/contexta(config-ca-crl)# cache-time n
```

where n is the number of minutes. For example, to specify that CRLs should be cached for seven hours, enter the following command.

```
hostname/contexta(config-ca-crl)# cache-time 420
```

Step 8 Configure whether the security appliance requires the NextUpdate field in CRLs. For more information about how the security appliance uses the NextUpdate field, see the [“About CRLs” section on page 39-3](#).

Do one of the following:

- To require the NextUpdate field, enter the **enforcenextupdate** command. This is the default setting.
- To allow the NextUpdate field to be absent in CRLs, enter the **no enforcenextupdate** command.

Step 9 If you specified LDAP as the retrieval protocol, perform the following steps:

a. Enter the following command to identify the LDAP server to the security appliance:

```
hostname/contexta(config-ca-crl)# ldap-defaults server
```

You can specify the server by DNS hostname or by IP address. You can also provide a port number if the server listens for LDAP queries on a port other than the default of 389. For example, the following command configures the security appliance to retrieve CRLs from an LDAP server whose hostname is ldap1.

```
hostname/contexta(config-ca-crl)# ldap-defaults ldap1
```



Note

If you use a hostname rather than an IP address to specify the LDAP server, be sure you have configured the security appliance to use DNS. For information about configuring DNS, see the **dns** commands in the *Cisco Security Appliance Command Reference*.

b. If LDAP server requires credentials to permit CRL retrieval, enter the following command:

```
hostname/contexta(config-ca-crl)# ldap-dn admin-DN password
```

For example:

```
hostname/contexta(config-ca-crl)# ldap-dn cn=admin,ou=devtest,o=engineering c001RunZ
```

Step 10 To test CRL configuration for the current trustpoint, use the **crypto ca crl request** command. This command retrieves the current CRL from the CA represented by the trustpoint you specify.

Step 11 Save the running configuration. Enter the **write memory** command.

Exporting and Importing Trustpoints

You can export and import keypairs and issued certificates associated with a trustpoint configuration. The security appliance supports PKCS12 format for the export and import of trustpoints.

This section includes the following topics:

- [Exporting a Trustpoint Configuration, page 39-15](#)
- [Importing a Trustpoint Configuration, page 39-15](#)

Exporting a Trustpoint Configuration

To export a trustpoint configuration with all associated keys and certificates in PKCS12 format, use the **crypto ca export** command. The security appliance displays the PKCS12 data in the terminal. You can copy the data. The trustpoint data is password protected; however, if you save the trustpoint data in a file, be sure the file is in a secure location.

The following example exports PKCS12 data for trustpoint Main using Wh0zits as the passphrase:

```
hostname (config)# crypto ca export Main pkcs12 Wh0zits

Exported pkcs12 follows:

[ PKCS12 data omitted ]

---End - This line not part of the pkcs12---

hostname (config)#
```

Importing a Trustpoint Configuration

To import the keypairs and issued certificates associated with a trustpoint configuration, use the **crypto ca import pkcs12** command in global configuration mode. The security appliance prompts you to paste the text to the terminal in base-64 format.

The key pair imported with the trustpoint is assigned a label matching the name of the trustpoint you create. For example, if an exported trustpoint used an RSA key labeled <Default-RSA-Key>, creating trustpoint named Main by importing the PKCS12 creates a key pair named Main, not <Default-RSA-Key>.



Note

If a security appliance has trustpoints that share the same CA, only one of the trustpoints sharing the CA can be used to validate user certificates. The **crypto ca import pkcs12** command can create this situation. Use the **support-user-cert-validation** command to control which trustpoint sharing a CA is used for validation of user certificates issued by that CA.

The following example manually imports PKCS12 data to the trustpoint Main with the passphrase Wh0zits:

```
hostname (config)# crypto ca import Main pkcs12 Wh0zits

Enter the base 64 encoded pkcs12.
End with a blank line or the word "quit" on a line by itself:
[ PKCS12 data omitted ]
quit
INFO: Import PKCS12 operation completed successfully
hostname (config)#
```

Configuring CA Certificate Map Rules

You can configure rules based on the Issuer and Subject fields of a certificate. Using the rules you create, you can map IPsec peer certificates to tunnel groups with the **tunnel-group-map** command. The security appliance supports one CA certificate map, which can contain many rules. For more information about using CA certificate map rules with tunnel groups, see the [“Creating a Certificate Group Matching Rule and Policy”](#) section on page 27-10.

To configure a CA certificate map rule, perform the following steps:

- Step 1** Enter CA certificate map configuration mode for the rule you want to configure. To do so, enter the **crypto ca certificate map command and specify the rule index number**. The following example enters CA certificate map mode for the rule with index number 1.

```
hostname(config)# crypto ca certificate map 1
hostname(config-ca-cert-map)#
```

- Step 2** Use the **issuer-name** and **subject-name** commands to configure the rule. These commands specify tests that the security appliance can apply to values found in the Issuer or Subject fields of certificates. The tests can apply to specific attributes or to the whole of the Issuer or Subject fields. You can configure many tests per rule, and all the tests you specify with these commands must be true for a rule to match a certificate. Valid operators in the **issuer-name** and **subject-name** commands are as follows.

Operator	Meaning
eq	The field or attribute must be identical to the value given.
ne	The field or attribute cannot be identical to the value given.
co	Part or all of the field or attribute must match the value given.
nc	No part of the field or attribute can match the value given.

For more information about the **issuer-name** and **subject-name** commands, see the *Cisco Security Appliance Command Reference*.

The following example specifies that any attribute within the Issuer field must contain the string cisco.

```
hostname(config-ca-cert-map)# issuer-name co cisco
hostname(config-ca-cert-map)#
```

The following example specifies that within the Subject field an Organizational Unit attribute must exactly match the string Engineering.

```
hostname(config-ca-cert-map)# subject-name attr ou eq Engineering
hostname(config-ca-cert-map)#
```

Map rules appear in the output of the **show running-config** command.

```
crypto ca certificate map 1
  issuer-name co cisco
  subject-name attr ou eq Engineering
```

- Step 3** When you have finished configuring the map rule, save your work. Enter the **write memory** command.



PART 4

System Administration



CHAPTER 40

Managing System Access

This chapter describes how to access the security appliance for system management through Telnet, SSH, and HTTPS. It also describes how to authenticate and authorize users and how to create login banners.

This chapter includes the following sections:

- [Allowing Telnet Access, page 40-1](#)
- [Allowing SSH Access, page 40-2](#)
- [Allowing HTTPS Access for ASDM, page 40-3](#)
- [Configuring ASDM and WebVPN on the Same Interface, page 40-4](#)
- [Configuring AAA for System Administrators, page 40-5](#)
- [Configuring a Login Banner, page 40-16](#)



Note

To access the security appliance interface for management access, you do not also need an access list allowing the host IP address. You only need to configure management access according to the sections in this chapter.

Allowing Telnet Access

The security appliance allows Telnet connections to the security appliance for management purposes. You cannot use Telnet to the lowest security interface unless you use Telnet inside an IPSec tunnel.

The security appliance allows a maximum of 5 concurrent Telnet connections per context, if available, with a maximum of 100 connections divided between all contexts.

To configure Telnet access to the security appliance, follow these steps:

- Step 1** To identify the IP addresses from which the security appliance accepts connections, enter the following command for each address or subnet:

```
hostname(config)# telnet source_IP_address mask source_interface
```

If there is only one interface, you can configure Telnet to access that interface as long as the interface has a security level of 100.

- Step 2** (Optional) To set the duration for how long a Telnet session can be idle before the security appliance disconnects the session, enter the following command:

```
hostname(config)# telnet timeout minutes
```

Set the timeout from 1 to 1440 minutes. The default is 5 minutes. The default duration is too short in most cases and should be increased until all pre-production testing and troubleshooting has been completed.

For example, to let a host on the inside interface with an address of 192.168.1.2 access the security appliance, enter the following command:

```
hostname(config)# telnet 192.168.1.2 255.255.255.255 inside
hostname(config)# telnet timeout 30
```

To allow all users on the 192.168.3.0 network to access the security appliance on the inside interface, enter the following command:

```
hostname(config)# telnet 192.168.3.0 255.255.255.0 inside
```

Allowing SSH Access

The security appliance allows SSH connections to the security appliance for management purposes. The security appliance allows a maximum of 5 concurrent SSH connections per context, if available, with a maximum of 100 connections divided between all contexts.

SSH is an application running on top of a reliable transport layer, such as TCP/IP, that provides strong authentication and encryption capabilities. The security appliance supports the SSH remote shell functionality provided in SSH Versions 1 and 2 and supports DES and 3DES ciphers.



Note

XML management over SSL and SSH are not supported.

This section includes the following topics:

- [Configuring SSH Access, page 40-2](#)
- [Using an SSH Client, page 40-3](#)

Configuring SSH Access

To configure SSH access to the security appliance, follow these steps:

- Step 1** To generate an RSA key pair, which is required for SSH, enter the following command:

```
hostname(config)# crypto key generate rsa modulus modulus_size
```

The modulus (in bits) is 512, 768, 1024, or 2048. The larger the key modulus size you specify, the longer it takes to generate an RSA. We recommend a value of 1024.

- Step 2** To save the RSA keys to persistent Flash memory, enter the following command:

```
hostname(config)# write mem
```

- Step 3** To identify the IP addresses from which the security appliance accepts connections, enter the following command for each address or subnet:

```
hostname(config)# ssh source_IP_address mask source_interface
```

The security appliance accepts SSH connections from all interfaces, including the one with the lowest security level.

- Step 4** (Optional) To set the duration for how long an SSH session can be idle before the security appliance disconnects the session, enter the following command:

```
hostname(config)# ssh timeout minutes
```

Set the timeout from 1 to 60 minutes. The default is 5 minutes. The default duration is too short in most cases and should be increased until all pre-production testing and troubleshooting has been completed.

For example, to generate RSA keys and let a host on the inside interface with an address of 192.168.1.2 access the security appliance, enter the following command:

```
hostname(config)# crypto key generate rsa modulus 1024
hostname(config)# write mem
hostname(config)# ssh 192.168.1.2 255.255.255.255 inside
hostname(config)# ssh 192.168.1.2 255.255.255.255 inside
hostname(config)# ssh timeout 30
```

To allow all users on the 192.168.3.0 network to access the security appliance on the inside interface, the following command:

```
hostname(config)# ssh 192.168.3.0 255.255.255.0 inside
```

By default SSH allows both version one and version two. To specify the version number enter the following command:

```
hostname(config)# ssh version version_number
```

The *version_number* can be 1 or 2.

Using an SSH Client

To gain access to the security appliance console using SSH, at the SSH client enter the username **pix** and enter the login password set by the **password** command (see the [“Changing the Login Password”](#) section on page 8-1).

When starting an SSH session, a dot (.) displays on the security appliance console before the SSH user authentication prompt appears, as follows:

```
hostname(config)# .
```

The display of the dot does not affect the functionality of SSH. The dot appears at the console when generating a server key or decrypting a message using private keys during SSH key exchange before user authentication occurs. These tasks can take up to two minutes or longer. The dot is a progress indicator that verifies that the security appliance is busy and has not hung.

Allowing HTTPS Access for ASDM

To use ASDM, you need to enable the HTTPS server, and allow HTTPS connections to the security appliance. All of these tasks are completed if you use the **setup** command. This section describes how to manually configure ASDM access.

The security appliance allows a maximum of 5 concurrent ASDM instances per context, if available, with a maximum of 32 ASDM instances between all contexts.

To configure ASDM access, follow these steps:

-
- Step 1** To identify the IP addresses from which the security appliance accepts HTTPS connections, enter the following command for each address or subnet:
- ```
hostname(config)# http source_IP_address mask source_interface
```
- Step 2** To enable the HTTPS server, enter the following command:
- ```
hostname(config)# http server enable
```
- Step 3** To specify the location of the ASDM image, enter the following command:
- ```
hostname(config)# asdm image disk0:/asdmfile
```
- 

For example, to enable the HTTPS server and let a host on the inside interface with an address of 192.168.1.2 access ASDM, enter the following commands:

```
hostname(config)# crypto key generate rsa modulus 1024
hostname(config)# write mem
hostname(config)# http server enable
hostname(config)# http 192.168.1.2 255.255.255.255 inside
```

To allow all users on the 192.168.3.0 network to access ASDM on the inside interface, enter the following command:

```
hostname(config)# http 192.168.3.0 255.255.255.0 inside
```

## Configuring ASDM and WebVPN on the Same Interface

The security appliance can support both WebVPN connections and HTTPS connections for ASDM administrative sessions simultaneously on the same interface. Both HTTPS and WebVPN use port 443 by default. Therefore, to enable both HTTPS and WebVPN on the same interface, you must specify a different port number for either HTTPS or WebVPN. An alternative is to configure WebVPN and HTTPS on different interfaces.

To specify a port for HTTPS, use the *port* argument of the **http server enable** command. The following example specifies that HTTPS ASDM sessions use port 444 on the outside interface. WebVPN is also enabled on the outside interface and uses the default port (443). With this configuration, remote users initiate ASDM sessions by entering `https://<outside_ip>:444` in the browser.

```
hostname(config)# http server enable 444
hostname(config)# http 192.168.3.0 255.255.255.0 outside
hostname(config)# webvpn
hostname(config-webvpn)# enable outside
```

To specify a port for WebVPN, use the **port** command from webvpn configuration mode. The next example enables WebVPN on port 444 of the outside interface. HTTPS for ASDM is also configured on the outside interface and uses the default port (443). With this configuration, remote users initiating WebVPN sessions enter `https://<outside_ip>:444` in the browser.

```
hostname(config)# http server enable
hostname(config)# http 192.168.3.0 255.255.255.0 outside
hostname(config)# webvpn
```

```
hostname(config-webvpn)# port 444
hostname(config-webvpn)# enable outside
```

## Configuring AAA for System Administrators

This section describes how to enable authentication and command authorization for system administrators. Before you configure AAA for system administrators, first configure the local database or AAA server according to [Chapter 13, “AAA Server and Local Database Support.”](#)

This section includes the following topics:

- [Configuring Authentication for CLI Access, page 40-5](#)
- [Configuring Authentication To Access Privileged EXEC Mode, page 40-6](#)
- [Configuring Command Authorization, page 40-7](#)
- [Configuring Command Accounting, page 40-14](#)
- [Viewing the Current Logged-In User, page 40-14](#)
- [Recovering from a Lockout, page 40-15](#)

## Configuring Authentication for CLI Access

If you enable CLI authentication, the security appliance prompts you for your username and password to log in. After you enter your information, you have access to user EXEC mode.

To enter privileged EXEC mode, enter the **enable** command or the **login** command (if you are using the local database only).

If you configure enable authentication (see the [“Configuring Authentication for the Enable Command” section on page 40-6](#)), the security appliance prompts you for your username and password. If you do not configure enable authentication, enter the system enable password when you enter the **enable** command (set by the **enable password** command). However, if you do not use enable authentication, after you enter the **enable** command, you are no longer logged in as a particular user. To maintain your username, use enable authentication.

For authentication using the local database, you can use the **login** command, which maintains the username but requires no configuration to turn on authentication.



### Note

Before the security appliance can authenticate a Telnet, SSH, or HTTP user, you must first configure access to the security appliance using the **telnet**, **ssh**, and **http** commands. These commands identify the IP addresses that are allowed to communicate with the security appliance.

To authenticate users who access the CLI, enter the following command:

```
hostname(config)# aaa authentication {telnet | ssh | http | serial} console {LOCAL |
server_group [LOCAL]}
```

The **http** keyword authenticates the ASDM client that accesses the security appliance using HTTPS. You only need to configure HTTP authentication if you want to use a AAA server. By default, ASDM uses the local database for authentication even if you do not configure this command. HTTP management authentication does not support the SDI protocol for a AAA server group.

If you use a AAA server group for authentication, you can configure the security appliance to use the local database as a fallback method if the AAA server is unavailable. Specify the server group name followed by **LOCAL** (**LOCAL** is case sensitive). We recommend that you use the same username and password in the local database as the AAA server because the security appliance prompt does not give any indication which method is being used.

You can alternatively use the local database as your main method of authentication (with no fallback) by entering **LOCAL** alone.

## Configuring Authentication To Access Privileged EXEC Mode

You can configure the security appliance to authenticate users with a AAA server or the local database when they enter the **enable** command. Alternatively, users are automatically authenticated with the local database when they enter the **login** command, which also accesses privileged EXEC mode depending on the user level in the local database.

This section includes the following topics:

- [Configuring Authentication for the Enable Command, page 40-6](#)
- [Authenticating Users Using the Login Command, page 40-6](#)

### Configuring Authentication for the Enable Command

You can configure the security appliance to authenticate users when they enter the **enable** command. If you do not authenticate the **enable** command, when you enter **enable**, the security appliance prompts for the system enable password (set by the **enable password** command), and you are no longer logged in as a particular user. Applying authentication to the **enable** command maintains the username. This feature is particularly useful when you perform command authorization, where usernames are important to determine the commands a user can enter.

To authenticate users who enter the **enable** command, enter the following command:

```
hostname(config)# aaa authentication enable console {LOCAL | server_group [LOCAL]}
```

The user is prompted for the username and password.

If you use a AAA server group for authentication, you can configure the security appliance to use the local database as a fallback method if the AAA server is unavailable. Specify the server group name followed by **LOCAL** (**LOCAL** is case sensitive). We recommend that you use the same username and password in the local database as the AAA server because the security appliance prompt does not give any indication which method is being used.

You can alternatively use the local database as your main method of authentication (with no fallback) by entering **LOCAL** alone.

### Authenticating Users Using the Login Command

From user EXEC mode, you can log in as any username in the local database using the **login** command.

This feature allows users to log in with their own username and password to access privileged EXEC mode, so you do not have to give out the system enable password to everyone. To allow users to access privileged EXEC mode (and all commands) when they log in, set the user privilege level to 2 (the default) through 15. If you configure local command authorization, then the user can only enter commands assigned to that privilege level or lower. See the [“Configuring Local Command Authorization” section on page 40-8](#) for more information.



**Caution**

If you add users to the local database who can gain access to the CLI and whom you do not want to enter privileged EXEC mode, you should configure command authorization. Without command authorization, users can access privileged EXEC mode (and all commands) at the CLI using their own password if their privilege level is 2 or greater (2 is the default). Alternatively, you can use a AAA server for authentication, or you can set all local users to level 1 so you can control who can use the system enable password to access privileged EXEC mode.

To log in as a user from the local database, enter the following command:

```
hostname> login
```

The security appliance prompts for your username and password. After you enter your password, the security appliance places you in the privilege level that the local database specifies.

## Configuring Command Authorization

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. If you want to control the access to commands, the security appliance lets you configure command authorization, where you can determine which commands that are available to a user.

This section includes the following topics:

- [Command Authorization Overview, page 40-7](#)
- [Configuring Local Command Authorization, page 40-8](#)
- [Configuring TACACS+ Command Authorization, page 40-11](#)

## Command Authorization Overview

You can use one of two command authorization methods:

- Local database—Configure the command privilege levels on the security appliance. When a local user authenticates with the **enable** command (or logs in with the **login** command), the security appliance places that user in the privilege level that is defined by the local database. The user can then access commands at the user's privilege level and below.

**Note**

You can use local command authorization without any users in the local database and without CLI or enable authentication. Instead, when you enter the **enable** command, you enter the system enable password, and the security appliance places you in level 15. You can then create enable passwords for every level, so that when you enter **enable n** (2 to 15), the security appliance places you in level *n*. These levels are not used unless you turn on local command authorization (see “[Configuring Local Command Authorization](#)” below). (See the *Cisco Security Appliance Command Reference* for more information about **enable**.)

- TACACS+ server—On the TACACS+ server, configure the commands that a user or group can use after they authenticate for CLI access. Every command that a user enters at the CLI is checked with the TACACS+ server.

## Configuring Local Command Authorization

Local command authorization places each user at a privilege level, and each user can enter any command at their privilege level or below. The security appliance 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.

This section includes the following topics:

- [Local Command Authorization Prerequisites, page 40-8](#)
- [Default Command Privilege Levels, page 40-8](#)
- [Assigning Privilege Levels to Commands and Enabling Authorization, page 40-9](#)
- [Viewing Command Privilege Levels, page 40-10](#)

### Local Command Authorization Prerequisites

Complete the following tasks as part of your command authorization configuration:

- Configure **enable** authentication. (See the [“Configuring Authentication To Access Privileged EXEC Mode” section on page 40-6](#).)

Alternatively, you can use the **login** command (which is the same as the **enable** command with authentication), which requires no configuration. We do not recommend this option because it is not as secure as enable authentication.

You can also use CLI authentication, but it is not required.

- Configure each user in the local database at a privilege level from 0 to 15.

### Default Command Privilege Levels

By default, the following commands are assigned to privilege level 0. All other commands are at level 15.

- **show checksum**
- **show curpriv**
- **enable** (enable mode)
- **help**
- **show history**
- **login**
- **logout**
- **pager**
- **show pager**
- **clear pager**
- **quit**
- **show version**

If you move any configure mode commands to a lower level than 15, be sure to move the **configure** command to that level as well, otherwise, the user will not be able to enter configuration mode.

To view all privilege levels, see the [“Viewing Command Privilege Levels” section on page 40-10](#).

## Assigning Privilege Levels to Commands and Enabling Authorization

To assign a command to a new privilege level, and enable authorization, follow these steps:

**Step 1** To assign a command to a privilege level, enter the following command:

```
hostname(config)# privilege [show | clear | cmd] level level [mode {enable | cmd}] command command
```

Repeat this command for each command you want to reassign.

See the following information about the options in this command:

- **show | clear | cmd**—These optional keywords let you set the privilege only for the show, clear, or configure form of the command. The configure form of the command is typically the form that causes a configuration change, either as the unmodified command (without the **show** or **clear** prefix) or as the **no** form. If you do not use one of these keywords, all forms of the command are affected.
- **level level**—A level between 0 and 15.
- **mode {enable | configure}**—If a command can be entered in user EXEC/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:
  - **enable**—Specifies both user EXEC mode and privileged EXEC mode.
  - **configure**—Specifies configuration mode, accessed using the **configure terminal** command.
- **command command**—The command you are configuring. You can only configure the privilege level of the *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.

Also, you cannot configure the privilege level of subcommands separately from the main command. For example, you can configure the **context** command, but not the **allocate-interface** command, which inherits the settings from the **context** command.

**Step 2** To enable local command authorization, enter the following command:

```
hostname(config)# aaa authorization command LOCAL
```

Even if you set command privilege levels, command authorization does not take place unless you enable command authorization with this command.

For example, the **filter** command has the following forms:

- **filter** (represented by the **configure** option)
- **show running-config filter**
- **clear configure filter**

You can set the privilege level separately for each form, or set the same privilege level for all forms by omitting this option. For example, set each form separately as follows.

```
hostname(config)# privilege show level 5 command filter
hostname(config)# privilege clear level 10 command filter
hostname(config)# privilege cmd level 10 command filter
```

Alternatively, you can set all filter commands to the same level:

```
hostname(config)# privilege level 5 command filter
```

The **show privilege** command separates the forms in the display.

The following example shows the use of the **mode** keyword. The **enable** command must be entered from user EXEC mode, while the **enable password** command, which is accessible in configuration mode, requires the highest privilege level.

```
hostname(config)# privilege cmd level 0 mode enable command enable
hostname(config)# privilege cmd level 15 mode cmd command enable
hostname(config)# privilege show level 15 mode cmd command enable
```

This example shows an additional command, the **configure** command, that uses the **mode** keyword:

```
hostname(config)# privilege show level 5 mode cmd command configure
hostname(config)# privilege clear level 15 mode cmd command configure
hostname(config)# privilege cmd level 15 mode cmd command configure
hostname(config)# privilege cmd level 15 mode enable command configure
```



**Note**

This last line is for the **configure terminal** command.

## Viewing Command Privilege Levels

The following commands let you view privilege levels for commands.

- To show all commands, enter the following command:

```
hostname(config)# show running-config all privilege all
```

- To show commands for a specific level, enter the following command:

```
hostname(config)# show running-config privilege level level
```

The *level* is an integer between 0 and 15.

- To show the level of a specific command, enter the following command:

```
hostname(config)# show running-config privilege command command
```

For example, for the **show running-config all privilege all** command, the system displays the current assignment of each CLI command to a privilege level. The following is sample output from the command.

```
hostname(config)# show running-config all privilege all
privilege show level 15 command aaa
privilege clear level 15 command aaa
privilege configure level 15 command aaa
privilege show level 15 command aaa-server
privilege clear level 15 command aaa-server
privilege configure level 15 command aaa-server
privilege show level 15 command access-group
privilege clear level 15 command access-group
privilege configure level 15 command access-group
privilege show level 15 command access-list
privilege clear level 15 command access-list
privilege configure level 15 command access-list
privilege show level 15 command activation-key
privilege configure level 15 command activation-key
....
```

The following command displays the command assignments for privilege level 10:

```
hostname(config)# show running-config privilege level 10
privilege show level 10 command aaa
```

The following command displays the command assignment for the **access-list** command:

```
hostname(config)# show running-config privilege command access-list
```

```
privilege show level 15 command access-list
privilege clear level 15 command access-list
privilege configure level 15 command access-list
```

## Configuring TACACS+ Command Authorization

If you enable TACACS+ command authorization, and a user enters a command at the CLI, the security appliance sends the command and username to the TACACS+ server to determine if the command is authorized.

When configuring command authorization with a TACACS+ server, do not save your configuration until you are sure it works the way you want. If you get locked out because of a mistake, you can usually recover access by restarting the security appliance. If you still get locked out, see the [“Recovering from a Lockout” section on page 40-15](#).

Be sure that your TACACS+ system is completely stable and reliable. The necessary level of reliability typically requires that you have a fully redundant TACACS+ server system and fully redundant connectivity to the security appliance. For example, in your TACACS+ server pool, include one server connected to interface 1, and another to interface 2. You can also configure local command authorization as a fallback method if the TACACS+ server is unavailable. In this case, you need to configure local users and command privilege levels according to the [“Configuring Command Authorization” section on page 40-7](#).

This section includes the following topics:

- [TACACS+ Command Authorization Prerequisites, page 40-11](#)
- [Configuring Commands on the TACACS+ Server, page 40-11](#)
- [Enabling TACACS+ Command Authorization, page 40-14](#)

### TACACS+ Command Authorization Prerequisites

Complete the following tasks as part of your command authorization configuration:

- Configure CLI authentication (see the [“Configuring Local Command Authorization” section on page 40-8](#)).
- Configure **enable** authentication (see the [“Configuring Authentication To Access Privileged EXEC Mode” section on page 40-6](#)).

### Configuring Commands on the TACACS+ Server

You can configure commands on a Cisco Secure Access Control Server (ACS) TACACS+ server as a shared profile component, for a group, or for individual users. For third-party TACACS+ servers, see your server documentation for more information about command authorization support.

See the following guidelines for configuring commands in Cisco Secure ACS Version 3.1; many of these guidelines also apply to third-party servers:

- The security appliance sends the commands to be authorized as “shell” commands, so configure the commands on the TACACS+ server as shell commands.

**Note**

Cisco Secure ACS might include a command type called “pix-shell.” Do not use this type for security appliance command authorization.

- The first word of the command is considered to be the main command. All additional words are considered to be arguments, which need to be preceded by **permit** or **deny**.

For example, to allow the **show running-configuration aaa-server** command, add **show running-configuration** to the command box, and type **permit aaa-server** in the arguments box.

- You can permit all arguments of a command that you do not explicitly deny by selecting the **Permit Unmatched Args** check box.

For example, you can configure just the **show** command, and then all the **show** commands are allowed. We recommend using this method so that you do not have to anticipate every variant of a command, including abbreviations and **?**, which shows CLI usage (see [Figure 40-1](#)).

**Figure 40-1** Permitting All Related Commands

- For commands that are a single word, you *must* permit unmatched arguments, even if there are no arguments for the command, for example **enable** or **help** (see [Figure 40-2](#)).

**Figure 40-2** 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 box, and **deny password** in the arguments box. Be sure to select the Permit Unmatched Args check box so that **enable** alone is still allowed (see [Figure 40-3](#)).

**Figure 40-3**      *Disallowing Arguments*

- When you abbreviate a command at the command line, the security appliance expands the prefix and main command to the full text, but it sends additional arguments to the TACACS+ server as you enter them.

For example, if you enter **sh log**, then the security appliance sends the entire command to the TACACS+ server, **show logging**. However, if you enter **sh log mess**, then the security appliance sends **show logging mess** to the TACACS+ server, and not the expanded command **show logging message**. You can configure multiple spellings of the same argument to anticipate abbreviations (see [Figure 40-4](#)).

**Figure 40-4**      *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**

## Enabling TACACS+ Command Authorization

Before you enable TACACS+ command authorization, be sure that you are logged into the security appliance as a user that is defined on the TACACS+ server, and that you have the necessary command authorization to continue configuring the security appliance. For example, you should log in as an admin user with all commands authorized. Otherwise, you could become unintentionally locked out.

To perform command authorization using a TACACS+ server, enter the following command:

```
hostname(config)# aaa authorization command tacacs+_server_group [LOCAL]
```

You can configure the security appliance to use the local database as a fallback method if the TACACS+ server is unavailable. To enable fallback, specify the server group name followed by **LOCAL** (**LOCAL** is case sensitive). We recommend that you use the same username and password in the local database as the TACACS+ server because the security appliance prompt does not give any indication which method is being used. Be sure to configure users in the local database (see the [“Configuring Command Authorization”](#) section on page 40-7) and command privilege levels (see the [“Configuring Local Command Authorization”](#) section on page 40-8).

## Configuring Command Accounting

You can send accounting messages to the TACACS+ accounting server when you enter any command other than **show** commands at the CLI. If you customize the command privilege level using the **privilege** command (see the [“Assigning Privilege Levels to Commands and Enabling Authorization”](#) section on page 40-9), you can limit which commands the security appliance accounts for by specifying a minimum privilege level. The security appliance does not account for commands that are below the minimum privilege level.

To enable command accounting, enter the following command:

```
hostname(config)# aaa accounting command [privilege level] server-tag
```

Where *level* is the minimum privilege level and *server-tag* is the name of the TACACS+ server group that to which the security appliance should send command accounting messages. The TACACS+ server group configuration must already exist. For information about configuring a AAA server group, see the [“Identifying AAA Server Groups and Servers”](#) section on page 13-12.

## Viewing the Current Logged-In User

To view the current logged-in user, enter the following command:

```
hostname# show curpriv
```

See the following sample **show curpriv** command output. A description of each field follows.

```
hostname# show curpriv
Username : admin
Current privilege level : 15
Current Mode/s : P_PRIV
```



Table 40-1 describes the **show curpriv** command output.

**Table 40-1** *show curpriv Display Description*

| Field                   | Description                                                                                                                                                                                                     |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Username                | Username. If you are logged in as the default user, the name is enable_1 (user EXEC) or enable_15 (privileged EXEC).                                                                                            |
| Current privilege level | Level from 0 to 15. Unless you configure local command authorization and assign commands to intermediate privilege levels, levels 0 and 15 are the only levels that are used.                                   |
| Current Mode/s          | Shows the access modes: <ul style="list-style-type: none"> <li>• P_UNPR—User EXEC mode (levels 0 and 1)</li> <li>• P_PRIV—Privileged EXEC mode (levels 2 to 15)</li> <li>• P_CONF—Configuration mode</li> </ul> |

## Recovering from a Lockout

In some circumstances, when you turn on command authorization or CLI authentication, you can be locked out of the security appliance CLI. You can usually recover access by restarting the security appliance. However, if you already saved your configuration, you might be locked out. Table 40-2 lists the common lockout conditions and how you might recover from them.

**Table 40-2** *CLI Authentication and Command Authorization Lockout Scenarios*

| Feature                                                                                  | Lockout Condition                                                             | Description                                                                                  | Workaround: Single Mode                                                                                                                                                                                                   | Workaround: Multiple Mode                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local CLI authentication                                                                 | No users in the local database                                                | If you have no users in the local database, you cannot log in, and you cannot add any users. | Log in and reset the passwords and <b>aaa</b> commands.                                                                                                                                                                   | Session into the security appliance from the switch. From the system execution space, you can change to the context and add a user.                                                                                                                                                                                                                                                                                                        |
| TACACS+ command authorization<br>TACACS+ CLI authentication<br>RADIUS CLI authentication | Server down or unreachable and you do not have the fallback method configured | If the server is unreachable, then you cannot log in or enter any commands.                  | <ol style="list-style-type: none"> <li>1. Log in and reset the passwords and AAA commands.</li> <li>2. Configure the local database as a fallback method so you do not get locked out when the server is down.</li> </ol> | <ol style="list-style-type: none"> <li>1. If the server is unreachable because the network configuration is incorrect on the security appliance, session into the security appliance from the switch. From the system execution space, you can change to the context and reconfigure your network settings.</li> <li>2. Configure the local database as a fallback method so you do not get locked out when the server is down.</li> </ol> |

**Table 40-2** *CLI Authentication and Command Authorization Lockout Scenarios (continued)*

| Feature                       | Lockout Condition                                                                      | Description                                                                                   | Workaround: Single Mode                                                                                                                                                                                                                            | Workaround: Multiple Mode                                                                                                                                                                                                                       |
|-------------------------------|----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TACACS+ command authorization | You are logged in as a user without enough privileges or as a user that does not exist | You enable command authorization, but then find that the user cannot enter any more commands. | Fix the TACACS+ server user account.<br><br>If you do not have access to the TACACS+ server and you need to configure the security appliance immediately, then log into the maintenance partition and reset the passwords and <b>aaa</b> commands. | Session into the security appliance from the switch. From the system execution space, you can change to the context and complete the configuration changes. You can also disable command authorization until you fix the TACACS+ configuration. |
| Local command authorization   | You are logged in as a user without enough privileges                                  | You enable command authorization, but then find that the user cannot enter any more commands. | Log in and reset the passwords and <b>aaa</b> commands.                                                                                                                                                                                            | Session into the security appliance from the switch. From the system execution space, you can change to the context and change the user level.                                                                                                  |

## Configuring a Login Banner

You can configure a message to display when a user connects to the security appliance, before a user logs in, or before a user enters privileged EXEC mode.

To configure a login banner, enter the following command in the system execution space or within a context:

```
hostname(config)# banner {exec | login | motd} text
```

Adds a banner to display at one of three times: when a user first connects (message-of-the-day (**motd**)), when a user logs in (**login**), and when a user accesses privileged EXEC mode (**exec**). When a user connects to the security appliance, the message-of-the-day banner appears first, followed by the login banner and prompts. After the user successfully logs in to the security appliance, the exec banner displays.

For the banner text, spaces are allowed but tabs cannot be entered using the CLI. You can dynamically add the hostname or domain name of the security appliance by including the strings **\$(hostname)** and **\$(domain)**. If you configure a banner in the system configuration, you can use that banner text within a context by using the **\$(system)** string in the context configuration.

To add more than one line, precede each line by the **banner** command.

For example, to add a message-of-the-day banner, enter:

```
hostname(config)# banner motd Welcome to $(hostname).
hostname(config)# banner motd Contact me at admin@example.com for any
hostname(config)# banner motd issues.
```



# CHAPTER 41

## Managing Software, Licenses, and Configurations

---

This chapter contains information about managing the security appliance software, licenses, and configurations, and includes the following sections:

- [Managing Licenses, page 41-1](#)
- [Viewing Files in Flash Memory, page 41-2](#)
- [Retrieving Files from Flash Memory, page 41-3](#)
- [Downloading Software or Configuration Files to Flash Memory, page 41-3](#)
- [Configuring the Application Image and ASDM Image to Boot, page 41-5](#)
- [Configuring the File to Boot as the Startup Configuration, page 41-6](#)
- [Performing Zero Downtime Upgrades for Failover Pairs, page 41-6](#)
- [Backing Up Configuration Files, page 41-8](#)
- [Configuring Auto Update Support, page 41-10](#)

## Managing Licenses

When you install the software, the existing activation key is extracted from the original image and stored in a file in the security appliance file system.

## Obtaining an Activation Key

To obtain an activation key, you will need a Product Authorization Key, which you can purchase from your Cisco account representative. After obtaining the Product Authorization Key, register it on the Web to obtain an activation key by performing the following steps:

---

**Step 1** Obtain the serial number for your security appliance by entering the following command:

```
hostname> show version | include Number
```

Enter the pipe character (|) as part of the command.

**Step 2** Connect a web browser to one of the following websites (the URLs are case-sensitive):

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

<http://www.cisco.com/go/license>

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

<http://www.cisco.com/go/license/public>

**Step 3** Enter the following information, when prompted:

- Your Product Authorization Key
- The serial number of your security appliance.
- Your email address.

The activation key will be automatically generated and sent to the email address that you provide.

## Entering a New Activation Key

To enter the activation key, enter the following command:

```
hostname(config)# activation-key key
```

The key is a four or five-element hexadecimal string with one space between each element. For example, a key in the correct form might look like the following key:

```
0xe02888da 0x4ba7bed6 0xf1c123ae 0xffd8624e
```

The leading 0x specifier is optional; all values are assumed to be hexadecimal.

If you are already in multiple context mode, enter this command in the system execution space.

Before entering the activation key, ensure that the image in Flash memory and the running image are the same. You can do this by rebooting the security appliance before entering the new activation key.



### Note

The activation key is not stored in your configuration file. The key is tied to the serial number of the device.

You must reboot the security appliance after entering the new activation key for the change to take effect in the running image.

This example shows how to change the activation key on the security appliance:

```
hostname(config)# activation-key 0xe02888da 0x4ba7bed6 0xf1c123ae 0xffd8624e
```

## Viewing Files in Flash Memory

You can view files in Flash memory and see information about the files.

- To view the files in Flash memory, enter the following command:

```
hostname# dir [flash: | disk0: | disk1:]
```

The **flash:** keyword represents the internal Flash memory on the PIX 500 series security appliance. You can enter **flash:** or **disk0:** for the internal Flash memory on the ASA 5500 series adaptive security appliance. The **disk1:** keyword represents the external Flash memory on the ASA. The internal Flash memory is the default.

For example:

```
hostname# dir

Directory of disk0:/
500 -rw- 4958208 22:56:20 Nov 29 2004 cdisk.bin
2513 -rw- 4634 19:32:48 Sep 17 2004 first-backup
2788 -rw- 21601 20:51:46 Nov 23 2004 backup.cfg
2927 -rw- 8670632 20:42:48 Dec 08 2004 asdmfile.bin
```

- To view extended information about a specific file, enter the following command:

```
hostname# show file information [path:/] filename
```

The default path is the root directory of the internal Flash memory (flash:/ or disk0:/).

For example:

```
hostname# show file information cdisk.bin

disk0:/cdisk.bin:
 type is image (XXX) []
 file size is 4976640 bytes version 7.0(1)
```

The file size listed is for example only.

## Retrieving Files from Flash Memory

You can retrieve files directly from the Flash disk by using an HTTPS connection with the following URL, in which you supply values for the ASA IP address and the filename:

```
https://ASA_IP/disk0/filename
```

This option is useful for customers who wish to do the following:

- Copy from the ASA binary image (as a backup).
- Copy from WebVPN capture files.
- Copy any other Flash files to a secure desktop.

## Downloading Software or Configuration Files to Flash Memory

You can download application images, ASDM images, configuration files, and other files to the internal Flash memory or, for the ASA 5500 series adaptive security appliance, to the external Flash memory from a TFTP, FTP, HTTP, or HTTPS server.

This section includes the following topics:

- [Downloading a File to a Specific Location, page 41-4](#)
- [Downloading a File to the Startup or Running Configuration, page 41-4](#)

## Downloading a File to a Specific Location

This section describes how to download the application image, ASDM software, a configuration file, or any other file that needs to be downloaded to Flash memory. To download a file to the running or startup configuration, see the [“Downloading a File to the Startup or Running Configuration” section on page 41-4](#).

For information about installing the Cisco SSL VPN client, see the [“Installing the SVC Software” section on page 38-2](#). For information about installing Cisco Secure Desktop on the security appliance, see the *Cisco Secure Desktop Configuration Guide for Cisco ASA 5500 Series Administrators*.

To configure the security appliance to use a specific application image or ASDM image if you have more than one installed, or have installed them in external Flash memory see the [“Configuring the Application Image and ASDM Image to Boot” section on page 41-5](#).



### Note

To successfully copy ASDM Version 5.0(5) to Flash memory, you must be running Version 7.0.

To configure the security appliance to use a specific configuration as the startup configuration, see the [“Configuring the File to Boot as the Startup Configuration” section on page 41-6](#).

For multiple context mode, you must be in the system execution space.

To download a file to Flash memory, see the following commands for each download server type:

- To copy from a TFTP server, enter the following command:

```
hostname# copy tftp://server[/path]/filename {flash:/ | disk0:/ |
disk1:/}[path/]filename
```

The **flash:/** keyword represents the internal Flash memory on the PIX 500 series security appliance. You can enter **flash:/** or **disk0:/** for the internal Flash memory on the ASA 5500 series adaptive security appliance. The **disk1:/** keyword represents the external Flash memory on the ASA.

- To copy from an FTP server, enter the following command:

```
hostname# copy ftp://[user[:password]@]server[/path]/filename {flash:/ | disk0:/ |
disk1:/}[path/]filename
```

- To copy from an HTTP or HTTPS server, enter the following command:

```
hostname# copy http[s]://[user[:password]@]server[:port]/[path]/filename {flash:/ |
disk0:/ | disk1:/}[path/]filename
```

- To use secure copy, first enable SSH, then enter the following command:

```
hostname# ssh scopy enable
```

Then from a Linux client enter the following command:

```
scp -v -pw password filename username@asa_address
```

The **-v** is for verbose, and if **-pw** is not specified you will be prompted for a password.

## Downloading a File to the Startup or Running Configuration

You can download a text file to the running or startup configuration from a TFTP, FTP, or HTTP(S) server, or from the Flash memory.

To copy a file to the startup configuration or running configuration, enter one of the following commands for the appropriate download server.



#### Note

When you copy a configuration to the running configuration, you merge the two configurations. A merge adds any new commands from the new configuration to the running configuration. If the configurations are the same, no changes occur. If commands conflict or if commands affect the running of the context, then the effect of the merge depends on the command. You might get errors, or you might have unexpected results.

- To copy from a TFTP server, enter the following command:

```
hostname# copy tftp://server[/path]/filename {startup-config | running-config}
```

- To copy from an FTP server, enter the following command:

```
hostname# copy ftp://[user[:password]@]server[/path]/filename {startup-config | running-config}
```

- To copy from an HTTP or HTTPS server, enter the following command:

```
hostname# copy http[s]://[user[:password]@]server[:port]/[path]/filename {startup-config | running-config}
```

- To copy from Flash memory, enter the following command:

```
hostname# copy {flash:/ | disk0:/ | disk1:/}[path/]filename {startup-config | running-config}
```

For example, to copy the configuration from a TFTP server, enter the following command:

```
hostname# copy tftp://209.165.200.226/configs/startup.cfg startup-config
```

To copy the configuration from an FTP server, enter the following command:

```
hostname# copy ftp://admin:letmein@209.165.200.227/configs/startup.cfg startup-config
```

To copy the configuration from an HTTP server, enter the following command:

```
hostname# copy http://209.165.200.228/configs/startup.cfg startup-config
```

## Configuring the Application Image and ASDM Image to Boot

By default, the security appliance boots the first application image it finds in internal Flash memory. It also boots the first ASDM image it finds in internal Flash memory, or if none exists there, then in external Flash memory. If you have more than one image, you should specify the image you want to boot. In the case of the ASDM image, if you do not specify the image to boot, even if you have only one image installed, then the security appliance inserts the **asdm image** command into the running configuration. To avoid problems with Auto Update (if configured), and to avoid the image search at each startup, you should specify the ASDM image you want to boot in the startup configuration.

- To configure the application image to boot, enter the following command:

```
hostname(config)# boot system url
```

where *url* is one of the following:

- {flash:/ | disk0:/ | disk1:/}[path/]filename

The **flash:/** keyword represents the internal Flash memory on the PIX 500 series security appliance. You can enter **flash:/** or **disk0:/** for the internal Flash memory on the ASA 5500 series adaptive security appliance. The **disk1:/** keyword represents the external Flash memory on the ASA.

- **tftp://[user[:password]@]server[:port]/[path/]filename**

This option is only supported for the ASA 5500 series adaptive security appliance.

You can enter up to four **boot system** command entries, to specify different images to boot from in order; the security appliance boots the first image it finds. Only one **boot system tftp:** command can be configured, and it must be the first one configured.

- To configure the ASDM image to boot, enter the following command:

```
hostname(config)# asdm image {flash:/ | disk0:/ | disk1:/}[path/]filename
```

## Configuring the File to Boot as the Startup Configuration

By default, the security appliance boots from a startup configuration that is a hidden file. You can alternatively set any configuration to be the startup configuration by entering the following command:

```
hostname(config)# boot config {flash:/ | disk0:/ | disk1:/}[path/]filename
```

The **flash:/** keyword represents the internal Flash memory on the PIX 500 series security appliance. You can enter **flash:/** or **disk0:/** for the internal Flash memory on the ASA 5500 series adaptive security appliance. The **disk1:/** keyword represents the external Flash memory on the ASA.

## Performing Zero Downtime Upgrades for Failover Pairs

The two units in a failover configuration should have the same major (first number) and minor (second number) software version. However, you do not need to maintain version parity on the units during the upgrade process; you can have different versions on the software running on each unit and still maintain failover support. To ensure long-term compatibility and stability, we recommend upgrading both units to the same version as soon as possible.

Table 41-1 shows the supported scenarios for performing zero-downtime upgrades on a failover pair.

**Table 41-1      Zero-Downtime Upgrade Support**

| Type of Upgrade     | Support                                                                                                                                                                                                                              |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Maintenance Release | <p>You can upgrade from any maintenance release to any other maintenance release within a minor release.</p> <p>For example, you can upgrade from 7.0(1) to 7.0(4) without first installing the maintenance releases in between.</p> |



**Table 41-1** Zero-Downtime Upgrade Support

| Type of Upgrade | Support                                                                                                                                                                                                                                                                     |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Minor Release   | <p>You can upgrade from a minor release to the next minor release. You cannot skip a minor release.</p> <p>For example, you can upgrade from 7.0 to 7.1. Upgrading from 7.0 directly to 7.2 is not supported for zero-downtime upgrades; you must first upgrade to 7.1.</p> |
| Major Release   | <p>You can upgrade from the last minor release of the previous version to the next major release.</p> <p>For example, you can upgrade from 7.9 to 8.0, assuming that 7.9 is the last minor version in the 7.x release.</p>                                                  |

For more details about upgrading the software on a failover pair, refer to the following topics:

- [Upgrading an Active/Standby Failover Configuration, page 41-7](#)
- [Upgrading and Active/Active Failover Configuration, page 41-8](#)

## Upgrading an Active/Standby Failover Configuration

To upgrade two units in an Active/Standby failover configuration, perform the following steps:

- 
- Step 1** Download the new software to both units, and specify the new image to load with the **boot system** command (see the “[Configuring the Application Image and ASDM Image to Boot](#)” section on [page 41-5](#)).
- Step 2** Reload the standby unit to boot the new image by entering the following command on the active unit:
- ```
active# failover reload-standby
```
- Step 3** When the standby unit has finished reloading, and is in the Standby Ready state, force the active unit to fail over to the standby unit by entering the following command on the active unit.



Note Use the **show failover** command to verify that the standby unit is in the Standby Ready state.

```
active# no failover active
```

- Step 4** Reload the former active unit (now the new standby unit) by entering the following command:
- ```
newstandby# reload
```
- Step 5** When the new standby unit has finished reloading, and is in the Standby Ready state, return the original active unit to active status by entering the following command:
- ```
newstandby# failover active
```
-

Upgrading and Active/Active Failover Configuration

To upgrade two units in an Active/Active failover configuration, perform the following steps:

-
- Step 1** Download the new software to both units, and specify the new image to load with the **boot system** command (see the “[Configuring the Application Image and ASDM Image to Boot](#)” section on [page 41-5](#)).
- Step 2** Make both failover groups active on the primary unit by entering the following command in the system execution space of the primary unit:
- ```
primary# failover active
```
- Step 3** Reload the secondary unit to boot the new image by entering the following command in the system execution space of the primary unit:
- ```
primary# failover reload-standby
```
- Step 4** When the secondary unit has finished reloading, and both failover groups are in the Standby Ready state on that unit, make both failover groups active on the secondary unit using the following command in the system execution space of the primary unit:



Note Use the **show failover** command to verify that both failover groups are in the Standby Ready state on the secondary unit.

```
primary# no failover active
```

- Step 5** Make sure both failover groups are in the Standby Ready state on the primary unit, and then reload the primary unit using the following command:
- ```
primary# reload
```
- Step 6** If the failover groups are configured with the **preempt** command, they will automatically become active on their designated unit after the preempt delay has passed. If the failover groups are not configured with the **preempt** command, you can return them to active status on their designated units using the **failover active group** command.
- 

## Backing Up Configuration Files

To back up your configuration, use one of the following methods:

- [Backing up the Single Mode Configuration or Multiple Mode System Configuration, page 41-9](#)
- [Backing Up a Context Configuration in Flash Memory, page 41-9](#)
- [Backing Up a Context Configuration within a Context, page 41-9](#)
- [Copying the Configuration from the Terminal Display, page 41-10](#)

## Backing up the Single Mode Configuration or Multiple Mode System Configuration

In single context mode or from the system configuration in multiple mode, you can copy the startup configuration or running configuration to an external server or to the local Flash memory:

- To copy to a TFTP server, enter the following command:

```
hostname# copy {startup-config | running-config} tftp://server[/path]/filename
```

- To copy to a FTP server, enter the following command:

```
hostname# copy {startup-config | running-config}
ftp://[user[:password]@]server[/path]/filename
```

- To copy to local Flash memory, enter the following command:

```
hostname# copy {startup-config | running-config} {flash:/ | disk0:/ |
disk1:/} [path/] filename
```

Be sure the destination directory exists. If it does not exist, first create the directory using the **mkdir** command.

## Backing Up a Context Configuration in Flash Memory

In multiple context mode, copy context configurations that are on the local Flash memory by entering one of the following commands in the system execution space:

- To copy to a TFTP server, enter the following command:

```
hostname# copy disk:[path/] filename tftp://server[/path]/filename
```

- To copy to a FTP server, enter the following command:

```
hostname# copy disk:[path/] filename ftp://[user[:password]@]server[/path]/filename
```

- To copy to local Flash memory, enter the following command:

```
hostname# copy {flash:/ | disk0:/ | disk1:/} [path/] filename {flash:/ | disk0:/ |
disk1:/} [path/] newfilename
```

Be sure the destination directory exists. If it does not exist, first create the directory using the **mkdir** command.

## Backing Up a Context Configuration within a Context

In multiple context mode, from within a context, you can perform the following backups:

- To copy the running configuration to the startup configuration server (connected to the admin context), enter the following command:

```
hostname/contexta# copy running-config startup-config
```

- To copy the running configuration to a TFTP server connected to the context network, enter the following command:

```
hostname/contexta# copy running-config tftp://server[/path]/filename
```

## Copying the Configuration from the Terminal Display

To print the configuration to the terminal, enter the following command:

```
hostname# show running-config
```

Copy the output from this command, then paste the configuration in to a text file.

## Configuring Auto Update Support

Auto Update is a protocol specification that allows an Auto Update server to download configurations and software images to many security appliances, and can provide basic monitoring of the security appliances from a central location.

The security appliance 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 security appliances configured as Auto Update clients.

**Note**

Auto Update is supported in single context mode only.

This section includes the following topics:

- [Configuring Communication with an Auto Update Server, page 41-10](#)
- [Configuring Client Updates as an Auto Update Server, page 41-12](#)
- [Viewing Auto Update Status, page 41-13](#)

## Configuring Communication with an Auto Update Server

To configure the security appliance as an Auto Update client, perform the following steps:

**Step 1** To specify the URL of the AUS, use the following command:

```
hostname(config)# auto-update server url [source interface] [verify-certificate]
```

Where *url* has the following syntax:

```
http[s]://[user:password@]server_ip[:port]/pathname
```

SSL is used when **https** is specified. The *user* and *password* arguments of the URL are used for Basic Authentication when logging in to the server. If you use the **write terminal**, **show configuration** or **show tech-support** commands to view the configuration, the user and password are replaced with '\*\*\*\*\*'.

The default port is 80 for HTTP and 443 for HTTPS.

The **source interface** argument specifies which interface to use when sending requests to the AUS. If you specify the same interface specified by the **management-access** command, the Auto Update requests travel over the same IPSec VPN tunnel used for management access.

The **verify-certificate** keyword verifies the certificate returned by the AUS.

**Step 2** (Optional) To identify the device ID to send when communicating with the AUS, enter the following command:



In the following example, a security appliance is configured to poll an AUS with IP address 209.165.200.224, at port number 1742, from the outside interface, with certificate verification.

It is also configured to use the hostname of the security appliance as the device ID. It is configured to poll every Friday and Saturday night at a random time between 10:00 p.m. and 11:00 p.m. On a failed polling attempt, it will try to reconnect to the AUS 10 times, and wait 3 minutes between attempts at reconnecting.

```
hostname(config)# auto-update server
https://jcrichon:farscape@209.165.200.224:1742/management source outside
verify-certificate
hostname(config)# auto-update device-id hostname
hostname(config)# auto-update poll-at Friday Saturday 22:00 randomize 60 2 10
```

## Configuring Client Updates as an Auto Update Server

The **client-update** command lets you enable the update for security appliances configured as Auto Update clients. It lets you specify the type of software component (asdm or boot image), the type or family of security appliance, revision numbers to which the update applies, and a URL or IP address from which to get the update.

To configure the security appliance as an Auto Update server, perform the following steps:

- 
- Step 1** In global configuration mode, enable client update by entering the command:

```
hostname(config)# client-update enable
hostname(config)#
```

- Step 2** Configure the parameters for the client update that you want to apply for the security appliances using the **client-update** command:

```
client-update {component {asdm | image} | device-id dev_string |
family family_name | type type} url url-string rev-nums rev-nums}
```

**component** {**asdm** | **image**} specifies the software component, either ASDM or the boot image of the security appliance.

**device-id** *dev\_string* specifies a unique string that the Auto Update client uses to identify itself. The maximum length is 63 characters.

**family** *family\_name* specifies the family name that the Auto Update client uses to identify itself. It can be asa, pix, or a text string with a maximum length of 7 characters.

**rev-nums** *rev-nums* specifies the software or firmware images for this client. Enter up to 4, in any order, separated by commas.

**type** *type* specifies the type of clients to notify of a client update. Because this command is also used to update Windows clients, the list of clients includes several Windows operating systems. The security appliances in the list include the following:

- pix-515: Cisco PIX 515 Firewall
- pix-515e: Cisco PIX 515E Firewall
- pix-525: Cisco PIX 525 Firewall
- pix-535: Cisco PIX 535 Firewall
- asa5505: Cisco 5505 Adaptive Security Appliance
- asa5510: Cisco 5510 Adaptive Security Appliance

- asa5520: Cisco 5520 Adaptive Security Appliance
- asa5540: Cisco Adaptive Security Appliance

**url** *url-string* specifies the URL for the software/firmware image. This URL must point to a file appropriate for this client. For all Auto Update clients, you must use the protocol “http://” or “https://” as the prefix for the URL.

Configure the parameters for the client update that you want to apply to all security appliances of a particular type. That is, specify the type of security appliance and the URL or IP address from which to get the updated image. In addition, you must specify a revision number. If the revision number of the remote security appliance matches one of the specified revision numbers, there is no need to update—the client ignores the update.

The following example configures a client update for Cisco 5520 Adaptive Security Appliances:

```
hostname(config)# client-update type asa5520 component asdm url
http://192.168.1.114/aus/asdm501.bin rev-nums 7.2(1)
```

---

## Viewing Auto Update Status

To view the Auto Update status, enter the following command:

```
hostname(config)# show auto-update
```

The following is sample output from the **show auto-update** command:

```
hostname(config)# show auto-update
Server: https://*****@209.165.200.224:1742/management.cgi?1276
Certificate will be verified
Poll period: 720 minutes, retry count: 2, retry period: 5 minutes
Timeout: none
Device ID: host name [corporate]
Next poll in 4.93 minutes
Last poll: 11:36:46 PST Tue Nov 13 2004
Last PDM update: 23:36:46 PST Tue Nov 12 2004
```







## CHAPTER 42

# Monitoring the Security Appliance

---

This chapter describes how to monitor the security appliance, and includes the following sections:

- [Using SNMP, page 42-1](#)
- [Configuring and Managing Logs, page 42-5](#)

## Using SNMP

This section describes how to use SNMP and includes the following topics:

- [SNMP Overview, page 42-1](#)
- [Enabling SNMP, page 42-3](#)

## SNMP Overview

The security appliance provides support for network monitoring using SNMP V1 and V2c. The security appliance supports traps and SNMP read access, but does not support SNMP write access.

You can configure the security appliance to send traps (event notifications) to a network management station (NMS), or you can use the NMS to browse the MIBs on the security appliance. MIBs are a collection of definitions, and the security appliance maintains a database of values for each definition. Browsing a MIB entails issuing an SNMP get request from the NMS. Use CiscoWorks for Windows or any other SNMP V1, MIB-II compliant browser to receive SNMP traps and browse a MIB.

[Table 42-1](#) lists supported MIBs and traps for the security appliance and, in multiple mode, for each context. You can download Cisco MIBs from the following website.

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

After you download the MIBs, compile them for your NMS.



### Note

In software versions 7.2(1), 8.0(2), and later, the SNMP information refreshes about every five seconds. As a result, we recommend that you wait for at least five seconds between consecutive polls.

**Table 42-1** *SNMP MIB and Trap Support*

| <b>MIB or Trap Support</b>      | <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SNMP core traps                 | <p>The security appliance sends the following core SNMP traps:</p> <ul style="list-style-type: none"> <li>• authentication—An SNMP request fails because the NMS did not authenticate with the correct community string.</li> <li>• linkup—An interface has transitioned to the “up” state.</li> <li>• linkdown—An interface is down, for example, if you removed the <b>nameif</b> command.</li> <li>• coldstart—The security appliance is running after a reload.</li> </ul> |
| MIB-II                          | <p>The security appliance supports browsing of the following groups and tables:</p> <ul style="list-style-type: none"> <li>• system</li> </ul>                                                                                                                                                                                                                                                                                                                                 |
| IF-MIB                          | <p>The security appliance supports browsing of the following tables:</p> <ul style="list-style-type: none"> <li>• ifTable</li> <li>• ifXTable</li> </ul>                                                                                                                                                                                                                                                                                                                       |
| RFC1213-MIB                     | <p>The security appliance supports browsing of the following table:</p> <ul style="list-style-type: none"> <li>• ip.ipAddrTable</li> </ul>                                                                                                                                                                                                                                                                                                                                     |
| SNMPv2-MIB                      | <p>The security appliance supports browsing the following:</p> <ul style="list-style-type: none"> <li>• snmp</li> </ul>                                                                                                                                                                                                                                                                                                                                                        |
| ENTITY-MIB                      | <p>The security appliance supports browsing of the following groups and tables:</p> <ul style="list-style-type: none"> <li>• entPhysicalTable</li> <li>• entLogicalTable</li> </ul> <p>The security appliance supports browsing of the following traps:</p> <ul style="list-style-type: none"> <li>• config-change</li> <li>• fru-insert</li> <li>• fru-remove</li> </ul>                                                                                                      |
| CISCO-IPSEC-FLOW-MONITOR-MIB    | <p>The security appliance supports browsing of the MIB.</p> <p>The security appliance supports browsing of the following traps:</p> <ul style="list-style-type: none"> <li>• start</li> <li>• stop</li> </ul>                                                                                                                                                                                                                                                                  |
| CISCO-REMOTE-ACCESS-MONITOR-MIB | <p>The security appliance supports browsing of the MIB.</p> <p>The security appliance supports browsing of the following traps:</p> <ul style="list-style-type: none"> <li>• session-threshold-exceeded</li> </ul>                                                                                                                                                                                                                                                             |
| CISCO-CRYPTO-ACCELERATOR-MIB    | <p>The security appliance supports browsing of the MIB.</p>                                                                                                                                                                                                                                                                                                                                                                                                                    |
| ALTIGA-GLOBAL-REG               | <p>The security appliance supports browsing of the MIB.</p>                                                                                                                                                                                                                                                                                                                                                                                                                    |

**Table 42-1** *SNMP MIB and Trap Support (continued)*

| MIB or Trap Support   | Description                                                                                                                                                                                                                                                                           |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cisco Firewall MIB    | The security appliance supports browsing of the following groups: <ul style="list-style-type: none"> <li>cfwSystem</li> </ul> The information is cfwSystem.cfwStatus, which relates to failover status, pertains to the entire device and not just a single context.                  |
| Cisco Memory Pool MIB | The security appliance supports browsing of the following table: <ul style="list-style-type: none"> <li>ciscoMemoryPoolTable—The memory usage described in this table applies only to the security appliance general-purpose processor, and not to the network processors.</li> </ul> |
| Cisco Process MIB     | The security appliance supports browsing of the following table: <ul style="list-style-type: none"> <li>cpmCPUTotalTable</li> </ul>                                                                                                                                                   |
| Cisco Syslog MIB      | The security appliance supports the following trap: <ul style="list-style-type: none"> <li>clogMessageGenerated</li> </ul> You cannot browse this MIB.                                                                                                                                |

## Enabling SNMP

The SNMP agent that runs on the security appliance performs two functions:

- Replies to SNMP requests from NMSs.
- Sends traps (event notifications) to NMSs.

To enable the SNMP agent and identify an NMS that can connect to the security appliance, follow these steps:

**Step 1** Ensure that the SNMP server on the security appliance is enabled by entering the following command:

```
hostname(config)# snmp-server enable
```

The SNMP server is enabled by default.

**Step 2** To identify the IP address of the NMS that can connect to the security appliance, enter the following command:

```
hostname(config)# snmp-server host interface_name ip_address [trap | poll] [community text] [version 1 | 2c] [udp-port port]
```

Specify **trap** or **poll** if you want to limit the NMS to receiving traps only or browsing (polling) only. By default, the NMS can use both functions.

SNMP traps are sent on UDP port 162 by default. You can change the port number using the **udp-port** keyword.

**Step 3** To specify the community string, enter the following command:

```
hostname(config)# snmp-server community key
```

The SNMP community string is a shared secret between the security appliance and the NMS. The key is a case-sensitive value up to 32 characters in length. Spaces are not permitted.

**Step 4** (Optional) To set the SNMP server location or contact information, enter the following command:

```
hostname(config)# snmp-server {contact | location} text
```

**Step 5** To enable the security appliance to send traps to the NMS, enter the following command:

```
hostname(config)# snmp-server enable traps [all | syslog | snmp [trap] [...] |
entity [trap] [...] | ipsec [trap] [...] | remote-access [trap]]
```

Enter this command for each feature type to enable individual traps or sets of traps, or enter the **all** keyword to enable all traps.

The default configuration has all **snmp** traps enabled (**snmp-server enable traps snmp authentication linkup linkdown coldstart**). You can disable these traps using the **no** form of this command with the **snmp** keyword. However, the **clear configure snmp-server** command restores the default enabling of SNMP traps.

If you enter this command and do not specify a trap type, then the default is **syslog**. (The default **snmp** traps continue to be enabled along with the **syslog** trap.)

Traps for **snmp** include:

- **authentication**
- **linkup**
- **linkdown**
- **coldstart**

Traps for **entity** include:

- **config-change**
- **fru-insert**
- **fru-remove**

Traps for **ipsec** include:

- **start**
- **stop**

Traps for remote-access include:

- **session-threshold-exceeded**

**Step 6** To enable system messages to be sent as traps to the NMS, enter the following command:

```
hostname(config)# logging history level
```

You must also enable **syslog** traps using the preceding **snmp-server enable traps** command.

**Step 7** To enable logging, so system messages are generated and can then be sent to an NMS, enter the following command:

```
hostname(config)# logging enable
```

---

The following example sets the security appliance to receive requests from host 192.168.3.2 on the inside interface.

```
hostname(config)# snmp-server host 192.168.3.2
hostname(config)# snmp-server location building 42
hostname(config)# snmp-server contact Pat lee
hostname(config)# snmp-server community ohwhatakeyisthee
```

# Configuring and Managing Logs

This section describes the logging functionality and configuration. It also describes the system log message format, options and variables.

- [Logging Overview, page 42-5](#)
- [Logging in Multiple Context Mode, page 42-5](#)
- [Enabling and Disabling Logging, page 42-6](#)
- [Configuring Log Output Destinations, page 42-7](#)
- [Filtering System Log Messages, page 42-14](#)
- [Customizing the Log Configuration, page 42-18](#)
- [Understanding System Log Messages, page 42-23](#)

## Logging Overview

The security appliance system logs provide you with information for monitoring and troubleshooting the security appliance. Using the logging feature, you can do the following:

- Specify which system log messages should be logged.
- Disable or change the severity level of a system log message.
- Specify one or more locations where system log 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 system log messages in groups, such as by severity level or class of message.
- Specify what happens to the contents of the internal buffer when the buffer becomes full: overwrite the buffer, send the buffer contents to an FTP server, or save the contents to internal Flash memory.

You can choose to send all system log messages, or subsets of system log messages, to any or all output locations. You can filter which system log messages are sent to which locations by the severity of the system log message, the class of the system log message, or by creating a custom log 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 that are related to the current context.

System messages that are generated in the system execution space, including failover messages, are viewed in the admin context along with messages generated in the admin context. You cannot configure logging or view any logging information in the system execution space.

You can configure the security appliance to include the context name with each message, which helps you differentiate context messages that are sent to a single syslog server. This feature also helps you to determine which messages are from the admin context and which are from the system; messages that originate in the system execution space use a device ID of **system**, and messages that originate in the admin context use the name of the admin context as the device ID. For more information about enabling logging device IDs, see the [“Including the Device ID in System Log Messages”](#) section on page 42-19.

## Enabling and Disabling Logging

This section describes how to enable and disable logging on the security appliance. It includes the following sections:

- [Enabling Logging to All Configured Output Destinations, page 42-6](#)
- [Disabling Logging to All Configured Output Destinations, page 42-6](#)
- [Viewing the Log Configuration, page 42-6](#)

### Enabling Logging to All Configured Output Destinations

The following command enables logging; however, you must also specify at least one output destination so that you can view or save the logged messages. If you do not specify an output destination, the security appliance does not save system log messages generated when events occur.

For more information about configuring log output destinations, see the [“Configuring Log Output Destinations” section on page 42-7](#).

To enable logging, enter the following command:

```
hostname(config)# logging enable
```

### Disabling Logging to All Configured Output Destinations

To disable all logging to all configured log output destinations, enter the following command:

```
hostname(config)# no logging enable
```

### Viewing the Log Configuration

To view the running log configuration, enter the following command:

```
hostname(config)# show logging
```

The output of the **show logging** command is similar to the following:

```
Syslog logging: enabled
 Facility: 16
 Timestamp logging: disabled
 Standby logging: disabled
 Deny Conn when Queue Full: disabled
 Console logging: disabled
 Monitor logging: disabled
 Buffer logging: disabled
 Trap logging: level errors, facility 16, 3607 messages logged
 Logging to infrastructure 10.1.2.3
 History logging: disabled
 Device ID: 'inside' interface IP address "10.1.1.1"
 Mail logging: disabled
 ASDM logging: disabled
```

## Configuring Log Output Destinations

This section describes how to specify where the security appliance should save or send the log messages it generates. To view logs generated by the security appliance, you must specify a log output destination. If you enable logging without specifying a log output destination, the security appliance generates messages but does not save them to a location from which you can view them.

This section includes the following topics:

- [Sending System Log Messages to a Syslog Server, page 42-7](#)
- [Sending System Log Messages to the Console Port, page 42-8](#)
- [Sending System Log Messages to an E-mail Address, page 42-9](#)
- [Sending System Log Messages to ASDM, page 42-10](#)
- [Sending System Log Messages to a Telnet or SSH Session, page 42-11](#)
- [Sending System Log Messages to the Log Buffer, page 42-12](#)

### Sending System Log Messages to a Syslog Server

This section describes how to configure the security appliance to send logs to a syslog server.

Configuring the security appliance to send logs to a syslog server enables you to archive logs, limited only by the available disk space on the server, and it enables you to manipulate log data after it is saved. For example, you could specify actions to be executed when certain types of system log 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.

The syslog server must run a program (known as a server) called syslogd. UNIX provides a syslog server as part of its operating system. For Windows 95 and Windows 98, obtain a syslogd server from another vendor.



#### Note

To start logging to a syslog server you define in this procedure, be sure to enable logging for all output locations. See the [“Enabling Logging to All Configured Output Destinations”](#) section on page 42-6. To disable logging, see the [“Disabling Logging to All Configured Output Destinations”](#) section on page 42-6.

To configure the security appliance to send system log messages to a syslog server, perform the following steps:

**Step 1** To designate a syslog server to receive the logs, enter the following command:

```
hostname(config)# logging host interface_name ip_address [tcp[/port] | udp[/port]]
[format emblem]
```

Where the **format emblem** keyword enables EMBLEM format logging for the syslog server. (UDP only).

The *interface\_name* argument specifies the interface through which you access the syslog server.

The *ip\_address* argument specifies the IP address of the syslog server.

The **tcp[/port]** or **udp[/port]** argument specifies that the security appliance should use TCP or UDP to send system log messages to the syslog server. The default protocol is UDP. You can configure the security appliance to send data to a syslog server using either UDP or TCP, but not both. If you specify TCP, the security appliance discovers when the syslog server fails and discontinues sending logs. If you

specify UDP, the security appliance continues to send logs regardless of whether the syslog server is operational. The *port* argument specifies the port that the syslog server listens to for system log messages. Valid port values are 1025 through 65535, for either protocol. The default UDP port is 514. The default TCP port is 1470.

For example:

```
hostname(config)# logging host dmz1 192.168.1.5
```

If you want to designate more than one syslog server as an output destination, enter a new command for each syslog server.

- Step 2** To specify which system log messages should be sent to the syslog server, enter the following command:

```
hostname(config)# logging trap {severity_level | message_list}
```

Where the *severity\_level* argument specifies the severity levels of messages to be sent to the syslog server. You can specify the severity level number (0 through 7) or name. For severity level names, see the “[Severity Levels](#)” section on page 42-23. For example, if you set the level to 3, then the security appliance sends system log messages for level 3, 2, 1, and 0.

The *message\_list* argument specifies a customized message list that identifies the system log messages to send to the syslog server. For information about creating custom message lists, see the “[Filtering System Log Messages with Custom Message Lists](#)” section on page 42-17.

The following example specifies that the security appliance should send to the syslog server all system log messages with a severity level of level 3 (errors) and higher. The security appliance will send messages with the severity level of 3, 2, and 1.

```
hostname(config)# logging trap errors
```

- Step 3** (Optional) If needed, set the logging facility to a value other than its default of 20 by entering the following command:

```
hostname(config)# logging facility number
```

Most UNIX systems expect the system log messages to arrive at facility 20.

- Step 4** (Optional) To continue to pass traffic when the TCP syslog server is down, enter the following command:

```
hostname(config)# logging host interface_name server_ip [tcp/port] [permit-hostdown]
```

## Sending System Log Messages to the Console Port

This section describes how to configure the security appliance to send logs to the console port.



### Note

To start logging to the console port as defined in this procedure, be sure to enable logging for all output locations. See the “[Enabling Logging to All Configured Output Destinations](#)” section on page 42-6. To disable logging, see the “[Disabling Logging to All Configured Output Destinations](#)” section on page 42-6.

To specify which system log messages should be sent to the console port, enter the following command:

```
hostname(config)# logging console {severity_level | message_list}
```



Where the *severity\_level* argument specifies the severity levels of messages to be sent to the console port. You can specify the severity level number (0 through 7) or name. For severity level names, see the [“Severity Levels” section on page 42-23](#). For example, if you set the level to 3, then the security appliance sends system log messages for level 3, 2, 1, and 0.

The *message\_list* argument specifies a customized message list that identifies the system log messages to send to the console port. For information about creating custom message lists, see the [“Filtering System Log Messages with Custom Message Lists” section on page 42-17](#).

The following example specifies that the security appliance should send to the syslog server all system log messages with a severity level of level 3 (errors) and higher. The security appliance will send messages with the severity of 3, 2, and 1.

```
hostname(config)# logging console errors
```

## Sending System Log Messages to an E-mail Address

You can configure the security appliance to send some or all system log messages to an e-mail address. When sent by e-mail, a system log message appears in the subject line of the e-mail message. For this reason, we recommend configuring this option to notify administrators of system log messages with high severity levels, such as critical, alert, and emergency.



### Note

To start logging to an e-mail address you define in this procedure, be sure to enable logging for all output locations. See the [“Enabling Logging to All Configured Output Destinations” section on page 42-6](#). To disable logging, see the [“Disabling Logging to All Configured Output Destinations” section on page 42-6](#).

To designate an e-mail address as an output destination, perform the following steps:

- Step 1** To specify the system log messages to be sent to one or more e-mail addresses, enter the following command:

```
hostname(config)# logging mail {severity_level | message_list}
```

Where the *severity\_level* argument specifies the severity levels of messages to be sent to the e-mail address. You can specify the severity level number (0 through 7) or name. For severity level names, see the [“Severity Levels” section on page 42-23](#). For example, if you set the level to 3, then the security appliance sends system log messages for level 3, 2, 1, and 0.

The *message\_list* argument specifies a customized message list that identifies the system log messages to send to the e-mail address. For information about creating custom message lists, see the [“Filtering System Log Messages with Custom Message Lists” section on page 42-17](#).

The following example uses a *message\_list* with the name “high-priority,” previously set up with the **logging list** command:

```
hostname(config)# logging mail high-priority
```

- Step 2** To specify the source e-mail address to be used when sending system log messages to an e-mail address, enter the following command:

```
hostname(config)# logging from-address email_address
```

For example:

```
hostname(config)# logging from-address xxx-001@example.com
```

- Step 3** Specify the recipient e-mail address to be used when sending system log messages to an e-mail destination. You can configure up to five recipient addresses. You must enter each recipient separately.

To specify a recipient address, enter the following command:

```
hostname(config)# logging recipient-address e-mail_address [severity_level]
```

If a severity level is not specified, the default severity level is used (error condition, severity level 3).

For example:

```
hostname(config)# logging recipient-address admin@example.com
```

- Step 4** To specify the SMTP server to be used when sending system log messages to an e-mail destination, enter the following command:

```
hostname(config)# smtp-server ip_address
```

For example:

```
hostname(config)# smtp-server 10.1.1.1
```

## Sending System Log Messages to ASDM

You can configure the security appliance to send system log messages to ASDM. The security appliance sets aside a buffer area for system log 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. For information about the internal log buffer, see the [“Sending System Log Messages to the Log Buffer” section on page 42-12](#).

When the ASDM log buffer is full, the security appliance deletes the oldest system log message to make room in the buffer for new system log messages. To control the number of system log messages retained in the ASDM log buffer, you can change the size of the buffer.

This section includes the following topics:

- [Configuring Logging for ASDM, page 42-10](#)
- [Clearing the ASDM Log Buffer, page 42-11](#)

### Configuring Logging for ASDM



#### Note

To start logging to ASDM as defined in this procedure, be sure to enable logging for all output locations. See the [“Enabling Logging to All Configured Output Destinations” section on page 42-6](#). To disable logging, see the [“Disabling Logging to All Configured Output Destinations” section on page 42-6](#).

To specify ASDM as an output destination, perform the following steps:

- Step 1** To specify which system log messages should go to ASDM, enter the following command:

```
hostname(config)# logging asdm {severity_level | message_list}
```

Where the *severity\_level* argument specifies the severity levels of messages to be sent to ASDM. You can specify the severity level number (0 through 7) or name. For severity level names, see the [“Severity Levels” section on page 42-23](#). For example, if you set the level to 3, then the security appliance sends system log messages for level 3, 2, 1, and 0.

The *message\_list* argument specifies a customized message list that identifies the system log messages to send to ASDM. For information about creating custom message lists, see the [“Filtering System Log Messages with Custom Message Lists”](#) section on page 42-17.

The following example shows how enable logging and send to the ASDM log buffer system log messages of severity levels 0, 1, and 2.

```
hostname(config)# logging asdm 2
```

- Step 2** To specify the number of system log messages retained in the ASDM log buffer, enter the following command:

```
hostname(config)# logging asdm-buffer-size num_of_msgs
```

Where *num\_of\_msgs* specifies the number of system log messages that the security appliance retains in the ASDM log buffer.

The following example shows how to set the ASDM log buffer size to 200 system log messages.

```
hostname(config)# logging asdm-buffer-size 200
```

## Clearing the ASDM Log Buffer

To erase the current contents of the ASDM log buffer, enter the following command:

```
hostname(config)# clear logging asdm
```

## Sending System Log Messages to a Telnet or SSH Session

Viewing system log messages in a Telnet or SSH session requires two steps:

1. Specify which messages should be sent to Telnet or SSH sessions.
2. View logs in the current session.

This section includes the following topics:

- [Configuring Logging for Telnet and SSH Sessions, page 42-11](#)
- [Viewing System Log Messages in the Current Session, page 42-12](#)

## Configuring Logging for Telnet and SSH Sessions



### Note

To start logging to a Telnet or SSH session as defined in this procedure, be sure to enable logging for all output locations. See the [“Enabling Logging to All Configured Output Destinations”](#) section on page 42-6. To disable logging, see the [“Disabling Logging to All Configured Output Destinations”](#) section on page 42-6.

To specify which messages should be sent to Telnet or SSH sessions, enter the following command:

```
hostname(config)# logging monitor {severity_level | message_list}
```

Where the *severity\_level* argument specifies the severity levels of messages to be sent to the session. You can specify the severity level number (0 through 7) or name. For severity level names, see the [“Severity Levels”](#) section on page 42-23. For example, if you set the level to 3, then the security appliance sends system log messages for level 3, 2, 1, and 0.

The *message\_list* argument specifies a customized message list that identifies the system log messages to send to the session. For information about creating custom message lists, see the [“Filtering System Log Messages with Custom Message Lists”](#) section on page 42-17.

## Viewing System Log Messages in the Current Session

- Step 1** After you log in to the security appliance, enable logging to the current session by entering the following command:

```
hostname# terminal monitor
```

This command enables logging only for the current session. If you log out, and then log in again, you need to reenter this command.

- Step 2** To disable logging to the current session, enter the following command:

```
hostname(config)# terminal no monitor
```

## Sending System Log Messages to the Log Buffer

If configured as an output destination, the log buffer serves as a temporary storage location for system log messages. New messages are appended to the end of the listing. When the buffer is full, that is, when the buffer wraps, old messages are overwritten as new messages are generated, unless you configure the security appliance to save the full buffer to another location.

This section includes the following topics:

- [Enabling the Log Buffer as an Output Destination, page 42-12](#)
- [Viewing the Log Buffer, page 42-13](#)
- [Automatically Saving the Full Log Buffer to Flash Memory, page 42-13](#)
- [Automatically Saving the Full Log Buffer to an FTP Server, page 42-14](#)
- [Saving the Current Contents of the Log Buffer to Internal Flash Memory, page 42-14](#)
- [Clearing the Contents of the Log Buffer, page 42-14](#)

## Enabling the Log Buffer as an Output Destination



### Note

To start logging to the buffer as defined in this procedure, be sure to enable logging for all output locations. See the [“Enabling Logging to All Configured Output Destinations”](#) section on page 42-6. To disable logging, see the [“Disabling Logging to All Configured Output Destinations”](#) section on page 42-6.

To enable the log buffer as a log output destination, enter the following command:

```
hostname(config)# logging buffered {severity_level | message_list}
```

Where the *severity\_level* argument specifies the severity levels of messages to be sent to the buffer. You can specify the severity level number (0 through 7) or name. For severity level names, see the [“Severity Levels”](#) section on page 42-23. For example, if you set the level to 3, then the security appliance sends system log messages for level 3, 2, 1, and 0.

The *message\_list* argument specifies a customized message list that identifies the system log messages to send to the buffer. For information about creating custom message lists, see the [“Filtering System Log Messages with Custom Message Lists”](#) section on page 42-17.

For example, to specify that messages with severity levels 1 and 2 should be saved in the log buffer, enter one of the following commands:

```
hostname(config)# logging buffered critical
```

or

```
hostname(config)# logging buffered level 2
```

For the *message\_list* option, specify the name of a message list containing criteria for selecting messages to be saved in the log buffer.

```
hostname(config)# logging buffered notif-list
```

## Viewing the Log Buffer

To view the log buffer, enter the following command:

```
hostname(config)# show logging
```

## Changing the Log Buffer Size

By default, the log buffer size is 4 KB. To change the size of the log buffer, enter the following command:

```
hostname(config)# logging buffer-size bytes
```

Where the *bytes* argument sets the amount of memory used for the log buffer, in bytes. For example, if you specify 8192, the security appliance uses 8 KB of memory for the log buffer.

The following example specifies that the security appliance uses 16 KB of memory for the log buffer:

```
hostname(config)# logging buffer-size 16384
```

## Automatically Saving the Full Log Buffer to Flash Memory

Unless configured otherwise, the security appliance address messages to the log buffer on a continuing basis, overwriting old messages when the buffer is full. If you want to keep a history of logs, you can configure the security appliance to send the buffer contents to another output location each time the buffer fills. Buffer contents can be saved either to internal Flash memory or to an FTP server.

When saving the buffer content to another location, the security appliance creates log files with names that use a default time-stamp format, as follows:

```
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.

While the security appliance writes the log buffer contents to internal Flash memory or an FTP server, it continues saving new messages to the log buffer.

To specify that messages in the log buffer should be saved to internal Flash memory each time the buffer wraps, enter the following command:

```
hostname(config)# logging flash-bufferwrap
```

## Automatically Saving the Full Log Buffer to an FTP Server

See the [“Saving the Current Contents of the Log Buffer to Internal Flash Memory”](#) section for more information about saving the buffer.

To specify that messages in the log buffer should be saved to an FTP server each time the buffer wraps, perform the following steps.

- Step 1** To enable the security appliance to send the log buffer contents to an FTP server every time the buffer wraps, enter the following command:

```
hostname(config)# logging ftp-bufferwrap
```

- Step 2** To identify the FTP server, entering the following command:

```
hostname(config)# logging ftp-server server path username password
```

Where the *server* argument specifies the IP address of the external FTP server

The *path* argument specifies the directory path on the FTP server where the log buffer data is to be saved. This path is relative to the FTP root directory.

The *username* argument specifies a username that is valid for logging into the FTP server.

The *password* argument specifies the password for the username specified.

For example:

```
hostname(config)# logging ftp-server 10.1.1.1 /syslogs logsupervisor 1luvMy10gs
```

## Saving the Current Contents of the Log Buffer to Internal Flash Memory

At any time, you can save the contents of the buffer to internal Flash memory. To save the current contents of the log buffer to internal Flash memory, enter the following command:

```
hostname(config)# logging savelog [savefile]
```

For example, the following example saves the contents of the log buffer to internal Flash memory using the file name latest-logfile.txt:

```
hostname(config)# logging savelog latest-logfile.txt
```

## Clearing the Contents of the Log Buffer

To erase the contents of the log buffer, enter the following command:

```
hostname(config)# clear logging buffer
```

## Filtering System Log Messages

This section describes how to specify which system log messages should go to output destinations, and includes the following topics:

- [Message Filtering Overview, page 42-15](#)
- [Filtering System Log Messages by Class, page 42-15](#)
- [Filtering System Log Messages with Custom Message Lists, page 42-17](#)

## Message Filtering Overview

You can filter generated system log messages so that only certain system log messages are sent to a particular output destination. For example, you could configure the security appliance to send all system log messages to one output destination and also to send a subset of those system log messages to a different output destination.

Specifically, you can configure the security appliance so that system log messages are directed to an output destination according to the following criteria:

- System log message ID number
- System log message severity level
- System log message class (equivalent to a functional area of the security appliance)

You customize the above criteria by creating a message list that you can specify when you set the output destination in the [“Configuring Log Output Destinations”](#) section on page 42-7.

You can alternatively configure the security appliance to send a particular message class to each type of output destination independently of the message list.

For example, you could configure the security appliance to send to the internal log buffer all system log messages with severity levels of 1, 2 and 3, send all system log messages in the “ha” class to a particular syslog server, or create a list of messages that you name “high-priority” that are sent to an e-mail address to notify system administrators of a possible problem.

## Filtering System Log Messages by Class

The system log message class provides a method of categorizing system log messages by type, equivalent to a feature or function of the security appliance. For example, the “vpnc” class denotes the VPN client.

This section includes the following topics:

- [Message Class Overview, page 42-15](#)
- [Sending All Messages in a Class to a Specified Output Destination, page 42-16](#)

### Message Class Overview

With logging classes, you can specify an output location for an entire category of system log messages with a single command.

You can use system log message classes in two ways:

- Issue the **logging class** command to specify an output location for an entire category of system log messages.
- Create a message list using the **logging list** command that specifies the message class. See the [“Filtering System Log Messages with Custom Message Lists”](#) section on page 42-17 for this method.

All system log messages in a particular class share the same initial 3 digits in their system log message ID numbers. For example, all system log message IDs that begin with the digits 611 are associated with the vpnc (VPN client) class. System log messages associated with the VPN client feature range from 611101 to 611323.

## Sending All Messages in a Class to a Specified Output Destination

When you configure all messages in a class to go to a type of output destination, this configuration overrides the configuration in the specific output destination command. For example, if you specify that messages at level 7 should go to the log buffer, and you also specify that ha class messages at level 3 should go to the buffer, then the latter configuration takes precedence.

To configure the security appliance to send an entire system log message class to a configured output destination, enter the following command:

```
hostname(config)# logging class message_class {buffered | console | history | mail |
monitor | trap} [severity_level]
```

Where the *message\_class* argument specifies a class of system log messages to be sent to the specified output destination. See [Table 42-2](#) for a list of system log message classes.

The **buffered**, **history**, **mail**, **monitor**, and **trap** keywords specify the output destination to which system log messages in this class should be sent. The **history** keyword enables SNMP logging. The **monitor** keyword enables Telnet and SSH logging. The **trap** keyword enables syslog server logging. Select one destination per command line entry. If you want to specify that a class should go to more than one destination, enter a new command for each output destination.

The *severity\_level* argument further restricts the system log messages to be sent to the output destination by specifying a severity level. For more information about message severity levels, see the [“Severity Levels”](#) section on [page 42-23](#).

The following example specifies that all system log messages related to the class ha (high availability, also known as failover) with a severity level of 1 (alerts) should be sent to the internal logging buffer.

```
hostname(config)# logging class ha buffered alerts
```

[Table 42-2](#) lists the system log message classes and the ranges of system log message IDs associated with each class.

**Table 42-2 System Log Message Classes and Associated Message ID Numbers**

| Class   | Definition                   | System Log Message ID Numbers                                                                                |
|---------|------------------------------|--------------------------------------------------------------------------------------------------------------|
| ha      | Failover (High Availability) | 101, 102, 103, 104, 210, 311, 709                                                                            |
| rip     | RIP Routing                  | 107, 312                                                                                                     |
| auth    | User Authentication          | 109, 113                                                                                                     |
| bridge  | Transparent Firewall         | 110, 220                                                                                                     |
| config  | Command interface            | 111, 112, 208, 308                                                                                           |
| sys     | System                       | 199, 211, 214, 216, 306, 307, 315, 414, 604, 605, 606, 610, 612, 614, 615, 701, 711                          |
| session | User Session                 | 106, 108, 201, 202, 204, 302, 303, 304, 305, 314, 405, 406, 407, 500, 502, 607, 608, 609, 616, 620, 703, 710 |
| ip      | IP Stack                     | 209, 215, 313, 317, 408                                                                                      |
| snmp    | SNMP                         | 212                                                                                                          |
| vpdn    | PPTP and L2TP Sessions       | 213, 403, 603                                                                                                |
| vpn     | IKE and IPsec                | 316, 320, 402, 404, 501, 602, 702, 713, 714, 715                                                             |
| ospf    | OSPF Routing                 | 318, 409, 503, 613                                                                                           |
| np      | Network Processor            | 319                                                                                                          |



**Table 42-2** System Log Message Classes and Associated Message ID Numbers

| Class (continued) | Definition                  | System Log Message ID Numbers |
|-------------------|-----------------------------|-------------------------------|
| <b>rm</b>         | Resource Manager            | 321                           |
| <b>ids</b>        | Intrusion Detection System  | 400, 401, 415                 |
| <b>vpnc</b>       | VPN Client                  | 611                           |
| <b>webvpn</b>     | Web-based VPN               | 716                           |
| <b>ca</b>         | PKI Certification Authority | 717                           |
| <b>e-mail</b>     | E-mail Proxy                | 719                           |
| <b>vpnlb</b>      | VPN Load Balancing          | 718                           |
| <b>vpnfo</b>      | VPN Failover                | 720                           |
| <b>npssl</b>      | NP SSL                      | 725                           |

## Filtering System Log Messages with Custom Message Lists

Creating a custom message list is a flexible way to exercise fine control over which system log messages are sent to which output destination. In a custom system log message list, you specify groups of system log messages using any or all of the following criteria: severity level, message IDs, ranges of system log message IDs, or by message class.

For example, message lists can be used to:

- Select system log messages with severity levels of 1 and 2 and send them to one or more e-mail addresses.
- Select all system log 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 criteria with a new command entry. It is possible to create a message list containing overlapping message selection criteria. If two criteria in a message list select the same message, the message is logged only once.

To create a customized list that the security appliance can use to select messages to be saved in the log buffer, perform the following steps:

**Step 1** Create a message list containing criteria for selecting messages by entering the following command:

```
hostname(config)# logging list name {level level [class message_class] |
message start_id[-end_id]}
```

Where the *name* argument specifies the name of the list. Do not use the names of severity levels as the name of a system log message list. Prohibited names include “emergencies,” “alert,” “critical,” “error,” “warning,” “notification,” “informational,” and “debugging.” Similarly, do not use the first three characters of these words at the beginning of a file name. For example, do not use a filename that starts with the characters “err.”

The **level level** argument specifies the severity level. You can specify the severity level number (0 through 7) or name. For severity level names, see the “Severity Levels” section on page 42-23. For example, if you set the level to 3, then the security appliance sends system log messages for level 3, 2, 1, and 0.

The **class message\_class** argument specifies a particular message class. See Table 42-2 on page 42-16 for a list of class names.

The **message** *start\_id*[-*end\_id*] argument specifies an individual system log message ID number or a range of numbers.

The following example creates a message list named `notif-list` that specifies messages with a severity level of 3 or higher should be saved in the log buffer:

```
hostname(config)# logging list notif-list level 3
```

**Step 2** (Optional) If you want to add more criteria for message selection to the list, enter the same command as in the previous step specifying the name of the existing message list and the additional criterion. Enter a new command for each criterion you want to add to the list.

The following example adds criteria to the message list: a range of message ID numbers, and the message class `ha` (high availability or failover).

```
hostname(config)# logging list notif-list 104024-105999
hostname(config)# logging list notif-list level critical
hostname(config)# logging list notif-list level warning class ha
```

The preceding example states that system log messages that match the criteria specified will be sent to the output destination. The specified criteria for system log messages to be included in the list are:

- System log message IDs that fall in the range of 104024 to 105999
- All system log messages with critical level or higher (emergency, alert, or critical)
- All `ha` class system log messages with warning level or higher (emergency, alert, critical, error, or warning)

A system log message is logged if it satisfies any of these conditions. If a system log satisfies more than one of the conditions, the message is logged only once.

## Customizing the Log Configuration

## Customizing the Log Configuration

This section describes other options for fine tuning the logging configuration. It includes the following topics:

- [Configuring the Logging Queue, page 42-19](#)
- [Including the Date and Time in System Log Messages, page 42-19](#)
- [Including the Device ID in System Log Messages, page 42-19](#)
- [Generating System Log Messages in EMBLEM Format, page 42-20](#)
- [Disabling a System Log Message, page 42-20](#)
- [Changing the Severity Level of a System Log Message, page 42-21](#)
- [Changing the Amount of Internal Flash Memory Available for Logs, page 42-22](#)

## Configuring the Logging Queue

The Cisco ASA has a fixed number of blocks in memory that can be allocated for buffering system log messages while they are waiting to be sent to the configured output destination. The number of blocks required depends on the length of the system log message queue and the number of syslog servers specified.

To specify the number of system log messages the security appliance can hold in its queue before sending them to the configured output destination, enter the following command:

```
hostname(config)# logging queue message_count
```

Where the *message\_count* variable specifies the number of system log messages that can remain in the system log message queue while awaiting processing. The default is 512 system log messages. A setting of 0 (zero) indicates unlimited system log messages, that is, the queue size is limited only by block memory availability.

To view the queue and queue statistics, enter the following command:

```
hostname(config)# show logging queue
```

## Including the Date and Time in System Log Messages

To specify that system log messages should include the date and time that the system log messages was generated, enter the following command:

```
hostname(config)# logging timestamp
```

## Including the Device ID in System Log Messages

To configure the security appliance to include a device ID in non-EMBLEM-format system log messages, enter the following command:

```
hostname(config)# logging device-id {context-name | hostname | ipaddress interface_name | string text}
```

You can specify only one type of device ID for the system log messages.

The **context-name** keyword indicates that the name of the current context should be used as the device ID (applies to multiple context mode only). If you enable the logging device ID for the admin context in multiple context mode, messages that originate in the system execution space use a device ID of **system**, and messages that originate in the admin context use the name of the admin context as the device ID.

The **hostname** keyword specifies that the hostname of the security appliance should be used as the device ID.

The **ipaddress interface\_name** argument specifies that the IP address of the interface specified as *interface\_name* should be used as the device ID. If you use the **ipaddress** keyword, the device ID becomes the specified security appliance interface IP address, regardless of the interface from which the system log message is sent. This keyword provides a single, consistent device ID for all system log messages that are sent from the device.

The **string text** argument specifies that the text string should be used as the device ID. The string can contain as many as 16 characters. You cannot use blank spaces or any of the following characters:

- & (ampersand)
- ' (single quote)

- " (double quote)
- < (less than)
- > (greater than)
- ? (question mark)

**Note**

If enabled, the device ID does not appear in EMBLEM-formatted system log messages or SNMP traps.

The following example enables the logging device ID for the FWSM:

```
hostname(config)# logging device-id hostname
```

The following example enables the logging device ID for a security context on the FWSM:

```
hostname(config)# logging device-id context-name
```

## Generating System Log Messages in EMBLEM Format

- To use the EMBLEM format for system log messages sent to destinations other than a syslog server, enter the following command:

```
hostname(config)# logging emblem
```

- To use the EMBLEM format for system log messages sent to a syslog server over UDP, specify the **format emblem** option when you configure the syslog server as a n output destination. See the [“Sending System Log Messages to a Syslog Server”](#) section on page 42-7 for more information about syslog servers. Enter the following command:

```
hostname(config)# logging host interface_name ip_address {tcp[/port] | udp[/port]}
[format emblem]
```

Where the *interface\_name* and *IP\_address* specifies the syslog server to receive the system log messages, **tcp[/port]** and **udp[/port]** indicate the protocol and port that should be used, and **format emblem** enables EMBLEM formatting for messages sent to the syslog server.

The Cisco ASA can send system log messages using either the UDP or TCP protocol; however, you can enable the EMBLEM format only for messages sent over UDP. The default protocol and port are UDP/514.

For example:

```
hostname(config)# logging host interface_1 122.243.006.123 udp format emblem
```

## Disabling a System Log Message

- To prevent the security appliance from generating a particular system log message, enter the following command:

```
hostname(config)# no logging message message_number
```

For example:

```
hostname(config)# no logging message 113019
```

- To reenable a disabled system log message, enter the following command:

```
hostname(config)# logging message message_number
```

For example:

```
hostname(config)# logging message 113019
```

- To see a list of disabled system log messages, enter the following command:

```
hostname(config)# show logging message
```

- To reenable logging of all disabled system log messages, enter the following command:

```
hostname(config)# clear config logging disabled
```

## Changing the Severity Level of a System Log Message

- To specify the logging level of a system log message, enter the following command:

```
hostname(config)# logging message message_ID level severity_level
```

The following example modifies the severity level of system log message ID 113019 from 4 (warnings) to 5 (notifications):

```
hostname(config)# logging message 113019 level 5
```

- To reset the logging level of a system log message to its default level, enter the following command.

```
hostname(config)# no logging message message_ID level current_severity_level
```

The following example modifies the severity level of system log message ID 113019 to its default value of 4 (warnings).

```
hostname(config)# no logging message 113019 level 5
```

- To see the severity level of a specific message, enter the following command:

```
hostname(config)# show logging message message_ID
```

- To see a list of system log messages with modified severity levels, enter the following command:

```
hostname(config)# show logging message
```

- To reset the severity level of all modified system log messages back to their defaults, enter the following command:

```
hostname(config)# clear configure logging level
```

The series of commands in the following example illustrates the use of the **logging message** command to control both whether a system log message is enabled and the severity level of the system log message.

```
hostname(config)# show logging message 403503
```

```
syslog 403503: default-level errors (enabled)
```

```
hostname(config)# logging message 403503 level 1
```

```
hostname(config)# show logging message 403503
```

```
syslog 403503: default-level errors, current-level alerts (enabled)
```

```
hostname(config)# no logging message 403503
```

```
hostname(config)# show logging message 403503
```

```
syslog 403503: default-level errors, current-level alerts (disabled)
```

```
hostname(config)# logging message 403503
```

```
hostname(config)# show logging message 403503
```

```
syslog 403503: default-level errors, current-level alerts (enabled)
```

```
hostname(config)# no logging message 403503 level 3
hostname(config)# show logging message 403503
syslog 403503: default-level errors (enabled)
```

## Changing the Amount of Internal Flash Memory Available for Logs

You can cause the security appliance to save the contents of the log buffer to Internal Flash memory in two ways:

- Configure logging so that the contents of the log buffer are saved to internal Flash memory each time the buffer wraps
- Enter a command instructing the security appliance to save the current contents of the log buffer to internal Flash memory immediately

By default, the security appliance can use up to 1 MB of internal Flash memory for log data. The default minimum amount of internal Flash memory that must be free for the security appliance 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 security appliance 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 are deleted, free memory would still be below the limit, the security appliance fails to save the new log file.

To modify the settings for the amount of internal Flash memory available for logs, perform the following steps:

- 
- Step 1** To specify the maximum amount of internal Flash memory available for saving log files, enter the following command:

```
hostname(config)# logging flash-maximum-allocation kbytes
```

Where *kbytes* specifies the maximum amount of internal Flash memory, in kilobytes, that can be used for saving log files.

The following example sets the maximum amount of internal Flash memory that can be used for log files to approximately 1.2 MB:

```
hostname(config)# logging flash-maximum-allocation 1200
```

- Step 2** To specify the minimum amount of internal Flash memory that must be free for the security appliance to save a log file, enter the following command:

```
hostname(config)# logging flash-minimum-free kbytes
```

Where *kbytes* specifies the minimum amount of internal Flash memory, in kilobytes, that must be available before the security appliance saves a new log file.

The following example specifies that the minimum amount of free internal Flash memory must be 4000 KB before the security appliance can save a new log file:

```
hostname(config)# logging flash-minimum-free 4000
```

---

## Understanding System Log Messages

This section describes the contents of system log messages generated by the security appliance. It includes the following topics:

- [System Log Message Format](#), page 42-23
- [Severity Levels](#), page 42-23

### System Log Message Format

System Log messages begin with a percent sign (%) and are structured as follows:

```
%PIX|ASA Level Message_number: Message_text
```

Field descriptions are as follows:

|                       |                                                                                                                                                                                                      |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>PIX ASA</i>        | Identifies the system log message facility code for messages generated by the Cisco ASA . This value is always PIX ASA .                                                                             |
| <i>Level</i>          | 1-7. The level reflects the severity of the condition described by the system log message. The lower the number, the more severe the condition. See <a href="#">Table 42-3</a> for more information. |
| <i>Message_number</i> | A unique 6-digit number that identifies the system log message.                                                                                                                                      |
| <i>Message_text</i>   | A text string describing the condition. This portion of the system log message sometimes includes IP addresses, port numbers, or usernames.                                                          |

### Severity Levels

[Table 42-3](#) lists the system log message severity levels.

**Table 42-3 System Log Message Severity Levels**

| Level Number | Level Keyword        | Description                       |
|--------------|----------------------|-----------------------------------|
| 0            | <b>emergencies</b>   | System unusable.                  |
| 1            | <b>alert</b>         | Immediate action needed.          |
| 2            | <b>critical</b>      | Critical condition.               |
| 3            | <b>error</b>         | Error condition.                  |
| 4            | <b>warning</b>       | Warning condition.                |
| 5            | <b>notification</b>  | Normal but significant condition. |
| 6            | <b>informational</b> | Informational message only.       |
| 7            | <b>debugging</b>     | Appears during debugging only.    |



#### Note

The security appliance does not generate system log messages with a severity level of 0 (emergencies). This level is provided in the **logging** command for compatibility with the UNIX system log feature, but is not used by the Cisco ASA .







## CHAPTER 43

# Troubleshooting the Security Appliance

---

This chapter describes how to troubleshoot the security appliance, and includes the following sections:

- [Testing Your Configuration, page 43-1](#)
- [Reloading the Security Appliance, page 43-6](#)
- [Performing Password Recovery, page 43-7](#)
- [Other Troubleshooting Tools, page 43-10](#)
- [Common Problems, page 43-11](#)

## Testing Your Configuration

This section describes how to test connectivity for the single mode security appliance or for each security context. The following steps describe how to ping the security appliance interfaces, and how to allow hosts on one interface to ping hosts on another interface.

We recommend that you only enable ping and debug messages during troubleshooting. When you are done testing the security appliance, follow the steps in the [“Disabling the Test Configuration” section on page 43-5](#).

This section includes the following topics:

- [Enabling ICMP Debug Messages and System Messages, page 43-1](#)
- [Pinging Security Appliance Interfaces, page 43-2](#)
- [Pinging Through the Security Appliance, page 43-4](#)
- [Disabling the Test Configuration, page 43-5](#)

## Enabling ICMP Debug Messages and System Messages

Debug messages and system messages can help you troubleshoot why your pings are not successful. The security appliance only shows ICMP debug messages for pings to the security appliance interfaces, and not for pings through the security appliance to other hosts. To enable debugging and system messages, perform the following steps:

- Step 1** To show ICMP packet information for pings to the security appliance interfaces, enter the following command:

```
hostname(config)# debug icmp trace
```

**Step 2** To set system messages to be sent to Telnet or SSH sessions, enter the following command:

```
hostname(config)# logging monitor debug
```

You can alternately use **logging buffer debug** to send messages to a buffer, and then view them later using the **show logging** command.

**Step 3** To send the system messages to your Telnet or SSH session, enter the following command:

```
hostname(config)# terminal monitor
```

**Step 4** To enable system messages, enter the following command:

```
hostname(config)# logging on
```

The following example shows a successful ping from an external host (209.165.201.2) to the security appliance outside interface (209.165.201.1):

```
hostname(config)# debug icmp trace
Inbound ICMP echo reply (len 32 id 1 seq 256) 209.165.201.1 > 209.165.201.2
Outbound ICMP echo request (len 32 id 1 seq 512) 209.165.201.2 > 209.165.201.1
Inbound ICMP echo reply (len 32 id 1 seq 512) 209.165.201.1 > 209.165.201.2
Outbound ICMP echo request (len 32 id 1 seq 768) 209.165.201.2 > 209.165.201.1
Inbound ICMP echo reply (len 32 id 1 seq 768) 209.165.201.1 > 209.165.201.2
Outbound ICMP echo request (len 32 id 1 seq 1024) 209.165.201.2 > 209.165.201.1
Inbound ICMP echo reply (len 32 id 1 seq 1024) 209.165.201.1 > 209.165.201.2
```

The preceding example shows the ICMP packet length (32 bytes), the ICMP packet identifier (1), and the ICMP sequence number (the ICMP sequence number starts at 0 and is incremented each time a request is sent).

## Pinging Security Appliance Interfaces

To test that the security appliance interfaces are up and running and that the security appliance and connected routers are routing correctly, you can ping the security appliance interfaces.



### Note

For security purposes the security appliance does not support far-end interface ping, that is pingging the IP address of the outside interface from the inside network.

To ping the security appliance interfaces, perform the following steps:


**Step 1** Create a sketch of your single mode security appliance or security context showing the interface names, security levels, and IP addresses.



### Note

Although this procedure uses IP addresses, the **ping** command also supports DNS names and names assigned to a local IP address with the **name** command.


The sketch should also include any directly connected routers, and a host on the other side of the router from which you will ping the security appliance. You will use this information for this procedure as well as the procedure in the [“Pinging Through the Security Appliance”](#) section on page 43-4. For example:

**Figure 43-1** Network Sketch with Interfaces, Routers, and Hosts

**Step 2** Ping each security appliance interface from the *directly connected* routers. For transparent mode, ping the management IP address.

This test ensures that the security appliance interfaces are active and that the interface configuration is correct.


A ping might fail if the security appliance interface is not active, the interface configuration is incorrect, or if a switch between the security appliance and router is down (see Figure 43-2). In this case, no debug messages or system messages appear on the security appliance, because the packet never reaches it.

**Figure 43-2** Ping Failure at Security Appliance Interface

If the ping reaches the security appliance, and the security appliance responds, you see debug messages like the following:

```
ICMP echo reply (len 32 id 1 seq 256) 209.165.201.1 > 209.165.201.2
ICMP echo request (len 32 id 1 seq 512) 209.165.201.2 > 209.165.201.1
```


If the ping reply does not return to the router, then you might have a switch loop or redundant IP addresses (see Figure 43-3).

**Figure 43-3 Ping Failure Because of IP Addressing Problems**

- Step 3** Ping each security appliance interface from a remote host. For transparent mode, ping the management IP address.

This test checks that the directly connected router can route the packet between the host and the security appliance, and that the security appliance can correctly route the packet back to the host.

A ping might fail if the security appliance does not have a route back to the host through the intermediate router (see Figure 43-4). In this case, the debug messages show that the ping was successful, but you see system message 110001 indicating a routing failure.

**Figure 43-4 Ping Failure Because the Security Appliance has no Route**

## Pinging Through the Security Appliance

After you successfully ping the security appliance interfaces, you should make sure traffic can pass successfully through the security appliance. For routed mode, this test shows that NAT is working correctly, if configured. For transparent mode, which does not use NAT, this test confirms that the security appliance is operating correctly; if the ping fails in transparent mode, contact Cisco TAC.

To ping between hosts on different interfaces, perform the following steps:

- Step 1** To add an access list allowing ICMP from any source host, enter the following command:

```
hostname(config)# access-list ICMPACL extended permit icmp any any
```

By default, when hosts access a lower security interface, all traffic is allowed through. However, to access a higher security interface, you need the preceding access list.

- Step 2** To assign the access list to each source interface, enter the following command:

```
hostname(config)# access-group ICMPACL in interface interface_name
```

Repeat this command for each source interface.

- Step 3** To enable the ICMP inspection engine, so ICMP responses are allowed back to the source host, enter the following commands:

```
hostname(config)# class-map ICMP-CLASS
```

```

hostname(config-cmap) # match access-list ICMPACL
hostname(config-cmap) # policy-map ICMP-POLICY
hostname(config-pmap) # class ICMP-CLASS
hostname(config-pmap-c) # inspect icmp
hostname(config-pmap-c) # service-policy ICMP-POLICY global

```

Alternatively, you can also apply the ICMPACL access list to the destination interface to allow ICMP traffic back through the security appliance.

**Step 4** Ping from the host or router through the source interface to another host or router on another interface. Repeat this step for as many interface pairs as you want to check.

If the ping succeeds, you see a system message confirming the address translation for routed mode (305009 or 305011) and that an ICMP connection was established (302020). You can also enter the **show xlate** and **show conns** commands to view this information.

If the ping fails for transparent mode, contact Cisco TAC.

For routed mode, the ping might fail because NAT is not configured correctly (see Figure 43-5). This is more likely if you enable NAT control. In this case, you see a system message showing that the NAT translation failed (305005 or 305006). If the ping is from an outside host to an inside host, and you do not have a static translation (which is required with NAT control), you see message 106010: deny inbound icmp.



**Note** The security appliance only shows ICMP debug messages for pings to the security appliance interfaces, and not for pings through the security appliance to other hosts.

**Figure 43-5** Ping Failure Because the Security Appliance is not Translating Addresses



## Disabling the Test Configuration

After you complete your testing, disable the test configuration that allows ICMP to and through the security appliance and that prints debug messages. If you leave this configuration in place, it can pose a serious security risk. Debug messages also slow the security appliance performance.

To disable the test configuration, perform the following steps:

- Step 1** To disable ICMP debug messages, enter the following command:
- ```
hostname(config) # no debug icmp trace
```
- Step 2** To disable logging, if desired, enter the following command:
- ```
hostname(config) # no logging on
```
- Step 3** To remove the ICMPACL access list, and also delete the related **access-group** commands, enter the following command:

```
hostname(config)# no access-list ICMPACL
```

**Step 4** (Optional) To disable the ICMP inspection engine, enter the following command:

```
hostname(config)# no service-policy ICMP-POLICY
```

---

## Traceroute

You can trace the route of a packet using the traceroute feature, which is accessed with the **traceroute** command. A traceroute works by sending UDP packets to a destination on an invalid port. Because the port is not valid, the routers along the way to the destination will respond with an ICMP Time Exceeded Message, and report that error back to the security appliance.

## Packet Tracer

In addition to capturing packets and the traceroute feature, it is possible to trace the lifespan of a packet through the security appliance to see if it is behaving as expected with the packet tracer tool. The packet tracer tool lets you do the following:

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

The **packet-tracer** command provides detailed information about the packets and how they are processed by the security appliance. In the instance that a command from the configuration did not cause the packet to drop, the **packet-tracer** command will provide information about the cause in an easily readable manner. For example if a packet was dropped because of an invalid header validation, a message is displayed that says, “packet dropped due to bad ip header (reason).”



### Note

The **packet-tracer** command can generate packets based on the 5 tuple information—source IP, destination IP, source port, destination port, and protocol. The packet tracer does not populate the data part of the packet and as a result some engine checks will not be applicable. The packet tracer will show that the packet is dropped not because it did not pass the inspection checks but because there is not enough data to test against the inspection checks. For example the packet tracer will show drops incorrectly for dns traffic if the dns inspection is enabled.

---

## Reloading the Security Appliance

In multiple mode, you can only reload from the system execution space. To reload the security appliance, enter the following command:

```
hostname# reload
```

# Performing Password Recovery

This section describes how to recover if you forget passwords, or you create a lockout situation because of AAA settings. You can also disable password recovery for extra security. This section includes the following topics:

- [Performing Password Recovery for the ASA 5500 Series Adaptive Security Appliance, page 43-7](#)
- [Password Recovery for the PIX 500 Series Security Appliance, page 43-8](#)
- [Disabling Password Recovery, page 43-9](#)
- [Resetting the Password on the SSM Hardware Module, page 43-10](#)

## Performing Password Recovery for the ASA 5500 Series Adaptive Security Appliance

To recover passwords, perform the following steps:

- 
- Step 1** Connect to the security appliance console port according to the [“Accessing the Command-Line Interface” section on page 2-4](#).
- Step 2** Power off the security appliance, and then power it on.
- Step 3** During the startup messages, press the **Escape** key when prompted to enter ROMMON.
- Step 4** To set the security appliance to ignore the startup configuration at reload, enter the following command:
- ```
rommon #1> confreg
```
- The security appliance displays the current configuration register value, and asks if you want to change the value:
- ```
Current Configuration Register: 0x00000011
Configuration Summary:
 boot TFTP image, boot default image from Flash on netboot failure
Do you wish to change this configuration? y/n [n]:
```
- Step 5** Record your current configuration register value, so you can restore it later.
- Step 6** At the prompt, enter **Y** to change the value.
- The security appliance prompts you for new values.
- Step 7** Accept the default values for all settings, except for the “disable system configuration?” value; at that prompt, enter **Y**.
- Step 8** Reload the security appliance by entering the following command:
- ```
rommon #2> boot
```
- The security appliance loads a default configuration instead of the startup configuration.
- Step 9** Enter privileged EXEC mode by entering the following command:
- ```
hostname> enable
```
- Step 10** When prompted for the password, press **Return**.
- The password is blank.
- Step 11** Load the startup configuration by entering the following command:

```
hostname# copy startup-config running-config
```

**Step 12** Enter global configuration mode by entering the following command:

```
hostname# configure terminal
```

**Step 13** Change the passwords in the configuration by entering the following commands, as necessary:

```
hostname(config)# password password
hostname(config)# enable password password
hostname(config)# username name password password
```

**Step 14** Change the configuration register to load the startup configuration at the next reload by entering the following command:

```
hostname(config)# config-register value
```

Where *value* is the configuration register value you noted in [Step 5](#) and 0x1 is the default configuration register. For more information about the configuration register, see the *Cisco Security Appliance Command Reference*.

**Step 15** Save the new passwords to the startup configuration by entering the following command:

```
hostname(config)# copy running-config startup-config
```

## Password Recovery for the PIX 500 Series Security Appliance

Performing password recovery on the security appliance erases the login password, enable password, and **aaa authentication console** commands. To erase these commands so you can log in with the default passwords, perform the following steps:

**Step 1** Download the PIX password tool from Cisco.com to a TFTP server accessible from the security appliance. See the link in the “Password Recovery Procedure for the PIX” document at the following URL:

```
http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_password_recovery09186a008009478b.shtml
```

**Step 2** Connect to the security appliance console port according to the [“Accessing the Command-Line Interface”](#) section on page 2-4.

**Step 3** Power off the security appliance, and then power it on.

**Step 4** Immediately after the startup messages appear, press the **Escape** key to enter monitor mode.

**Step 5** Configure the network settings for the interface that accesses the TFTP server by entering the following commands:

```
monitor> interface interface_id
monitor> address interface_ip
monitor> server tftp_ip
monitor> file pw_tool_name
monitor> gateway gateway_ip
```

**Step 6** Download the PIX password tool from the TFTP server by entering the following command:

```
monitor> tftp
```

If you have trouble reaching the server, you can enter the **ping address** command to test the connection.



**Step 7** At the “Do you wish to erase the passwords?” prompt, enter **Y**.

You can now log in with the default login password of “cisco” and the blank enable password.

The following example shows the PIX password recovery with the TFTP server on the outside interface:

```
monitor> interface 0
0: i8255X @ PCI(bus:0 dev:13 irq:10)
1: i8255X @ PCI(bus:0 dev:14 irq:7)

Using 0: i82559 @ PCI(bus:0 dev:13 irq:10), MAC: 0050.54ff.82b9
monitor> address 10.21.1.99
address 10.21.1.99
monitor> server 172.18.125.3
server 172.18.125.3
monitor> file np70.bin
file np52.bin
monitor> gateway 10.21.1.1
gateway 10.21.1.1
monitor> ping 172.18.125.3
Sending 5, 100-byte 0xf8d3 ICMP Echoes to 172.18.125.3, timeout is 4 seconds:
!!!!
Success rate is 100 percent (5/5)
monitor> tftp
tftp np52.bin@172.18.125.3 via 10.21.1.1.....
Received 73728 bytes

Cisco PIX password tool (4.0) #0: Tue Aug 22 23:22:19 PDT 2005
Flash=i28F640J5 @ 0x300
BIOS Flash=AT29C257 @ 0xd8000

Do you wish to erase the passwords? [yn] y
Passwords have been erased.

Rebooting....
```

## Disabling Password Recovery

You might want to disable password recovery to ensure that unauthorized users cannot use the password recovery mechanism to compromise the security appliance. To disable password recovery, enter the following command:

```
hostname(config)# no service password-recovery
```

On the ASA 5500 series adaptive security appliance, the **no service password-recovery** command prevents a user from entering ROMMON with the configuration intact. When a user enters ROMMON, the security appliance prompts the user to erase all Flash file systems. The user cannot enter ROMMON without first performing this erasure. If a user chooses not to erase the Flash file system, the security appliance reloads. Because password recovery depends on using ROMMON 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 recover 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 informational purposes only; when you enter the command at the CLI prompt, the setting is saved in NVRAM. The only way to change the setting is to enter the command at the CLI prompt. Loading a new configuration with a different version of the command does not change the setting. If you disable

password recovery when the security appliance is configured to ignore the startup configuration at startup (in preparation for password recovery), then the security appliance changes the setting to boot the startup configuration as usual. If you use failover, and the standby unit is configured to ignore the startup configuration, then the same change is made to the configuration register when the **no service password recovery** command replicates to the standby unit.

On the PIX 500 series security appliance, the **no service password-recovery** command forces the PIX password tool to prompt the user to erase all Flash file systems. The user cannot use the PIX password tool without first performing this erasure. If a user chooses not to erase the Flash file system, the security appliance reloads. Because password recovery depends on maintaining the existing configuration, this erasure prevents you from recovering a password. However, disabling password recovery prevents unauthorized users from viewing the configuration or inserting different passwords. In this case, to recover the system to an operating state, load a new image and a backup configuration file, if available.

## Resetting the Password on the SSM Hardware Module

To reset the password to the default of “cisco” on the SSM hardware module, perform the following steps:

---

**Step 1** Make sure the SSM hardware module is in the Up state and supports password reset.

**Step 2** Enter the following command:

```
hostname (config)# hw-module module 1 password-reset
```

Where *1* is the specified slot number on the SSM hardware module.



**Note**

On the AIP SSM, entering this command reboots the hardware module. The module is offline until the rebooting is finished. Enter the **show module** command to monitor the module status. The AIP SSM supports this command in version 6.0 and later.

On the CSC SSM, entering this command resets web services on the hardware module after the password has been reset. You may lose connection to ASDM or be logged out of the hardware module. The CSC SSM supports this command in the most recent version of 6.1, dated November 2006.

---

```
Reset the password on module in slot 1? [confirm]
```

**Step 3** Enter **y** to confirm.

---

## Other Troubleshooting Tools

The security appliance provides other troubleshooting tools to be used in conjunction with Cisco TAC:

- [Viewing Debug Messages, page 43-11](#)
- [Capturing Packets, page 43-11](#)
- [Viewing the Crash Dump, page 43-11](#)

## Viewing Debug Messages

Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use **debug** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco TAC. Moreover, it is best to use **debug** commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that increased debug command processing overhead will affect system use. To enable debug messages, see the **debug** commands in the *Cisco Security Appliance Command Reference*.

## Capturing Packets

Capturing packets is sometimes useful when troubleshooting connectivity problems or monitoring suspicious activity. We recommend contacting Cisco TAC if you want to use the packet capture feature. See the **capture** command in the *Cisco Security Appliance Command Reference*.

## Viewing the Crash Dump

If the security appliance crashes, you can view the crash dump information. We recommend contacting Cisco TAC if you want to interpret the crash dump. See the **show crashdump** command in the *Cisco Security Appliance Command Reference*.

## Common Problems

This section describes common problems with the security appliance, and how you might resolve them.

**Symptom** The context configuration was not saved, and was lost when you reloaded.

**Possible Cause** You did not save each context within the context execution space. If you are configuring contexts at the command line, you did not save the context before you changed to the next context.

**Recommended Action** Save each context within the context execution space using the **copy run start** command. You cannot save contexts from the system execution space.

**Symptom** You cannot make a Telnet connection or SSH to the security appliance interface.

**Possible Cause** You did not enable Telnet or SSH to the security appliance.

**Recommended Action** Enable Telnet or SSH to the security appliance according to the [“Allowing Telnet Access”](#) section on page 40-1 or the [“Allowing SSH Access”](#) section on page 40-2.

**Symptom** You cannot ping the security appliance interface.

**Possible Cause** You disabled ICMP to the security appliance.

**Recommended Action** Enable ICMP to the security appliance for your IP address using the **icmp** command.

**Symptom** You cannot ping through the security appliance, even though the access list allows it.

**Possible Cause** You did not enable the ICMP inspection engine or apply access lists on both the ingress and egress interfaces.

**Recommended Action** Because ICMP is a connectionless protocol, the security appliance does not automatically allow returning traffic through. In addition to an access list on the ingress interface, you either need to apply an access list to egress interface to allow replying traffic, or enable the ICMP inspection engine, which treats ICMP connections as stateful connections.

**Symptom** Traffic does not pass between two interfaces on the same security level.

**Possible Cause** You did not enable the feature that allows traffic to pass between interfaces on the same security level.

**Recommended Action** Enable this feature according to the [“Allowing Communication Between Interfaces on the Same Security Level”](#) section on page 7-6.



## **PART 5**

### **Reference**







# APPENDIX **A**

## Feature Licenses and Specifications

---

This appendix describes the feature licenses and specifications. This appendix includes the following sections:

- [Supported Platforms and Feature Licenses, page A-1](#)
- [Security Services Module Support, page A-9](#)
- [VPN Specifications, page A-10](#)

## Supported Platforms and Feature Licenses

This software version supports the following platforms; see the associated tables for the feature support for each model:

- ASA 5505, [Table A-1](#)
- ASA 5510, [Table A-2](#)
- ASA 5520, [Table A-3](#)
- ASA 5540, [Table A-4](#)
- ASA 5550, [Table A-5](#)
- PIX 515/515E, [Table A-6](#)
- PIX 525, [Table A-7](#)
- PIX 535, [Table A-8](#)



### Note

Items that are in italics are separate, optional licenses that you can replace the base license. You can mix and match licenses, for example, the 10 security context license plus the Strong Encryption license; or the 500 WebVPN license plus the GTP/GPRS license; or all four licenses together.

---

**Table A-1 ASA 5505 Adaptive Security Appliance License Features**

| ASA 5505                               | Base License                                                                          |                                            | Security Plus                         |                                            |
|----------------------------------------|---------------------------------------------------------------------------------------|--------------------------------------------|---------------------------------------|--------------------------------------------|
| Users, concurrent <sup>1</sup>         | 10                                                                                    | <i>Optional Licenses:</i>                  | 10                                    | <i>Optional Licenses:</i>                  |
|                                        |                                                                                       | 50    Unlimited                            |                                       | 50    Unlimited                            |
| Security Contexts                      | No support                                                                            |                                            | No support                            |                                            |
| VPN Sessions <sup>2</sup>              | 10 combined IPSec and WebVPN                                                          |                                            | 25 combined IPSec and WebVPN          |                                            |
| Max. IPSec Sessions                    | 10                                                                                    |                                            | 25                                    |                                            |
| Max. WebVPN Sessions                   | 2                                                                                     | <i>Optional License: 10</i>                | 2                                     | <i>Optional License: 10</i>                |
| VPN Load Balancing                     | No support                                                                            |                                            | No support                            |                                            |
| Failover                               | None                                                                                  |                                            | Active/Standby (no Stateful Failover) |                                            |
| GTP/GPRS                               | No support                                                                            |                                            | No support                            |                                            |
| Maximum VLANs/Zones                    | 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone) |                                            | 20                                    |                                            |
| Maximum VLAN Trunks                    | No support                                                                            |                                            | Unlimited                             |                                            |
| Concurrent Firewall Conns <sup>3</sup> | 10 K                                                                                  |                                            | 25 K                                  |                                            |
| Max. Physical Interfaces               | Unlimited, assigned to VLANs/zones                                                    |                                            | Unlimited, assigned to VLANs/zones    |                                            |
| Encryption                             | Base (DES)                                                                            | <i>Optional license: Strong (3DES/AES)</i> | Base (DES)                            | <i>Optional license: Strong (3DES/AES)</i> |
| Minimum RAM                            | 128 MB                                                                                |                                            | 128 MB                                |                                            |

1. In routed mode, hosts on the inside (Business and Home VLANs) count towards the limit only when they communicate with the outside (Internet VLAN). Internet hosts are not counted towards the limit. Hosts that initiate traffic between Business and Home are also not counted towards the limit. The interface associated with the default route is considered to be the Internet interface. If there is no default route, hosts on all interfaces are counted toward the limit. In transparent mode, the interface with the lowest number of hosts is counted towards the host limit. See the **show local-host command** to view the host limits.
2. Although the maximum IPSec and WebVPN sessions add up to more than the maximum VPN sessions, the combined sessions should not exceed the VPN session limit. If you exceed the maximum VPN sessions, you can overload the security appliance, so be sure to size your network appropriately.
3. The concurrent firewall connections are based on a traffic mix of 80% TCP and 20% UDP, with one host and one dynamic translation for every four connections.



**Table A-2 ASA 5510 Adaptive Security Appliance License Features**

| ASA 5510                               | Base License                                                          |                           |                                            |    |     |     | Security Plus                   |                           |                                            |    |     |     |
|----------------------------------------|-----------------------------------------------------------------------|---------------------------|--------------------------------------------|----|-----|-----|---------------------------------|---------------------------|--------------------------------------------|----|-----|-----|
| Users, concurrent                      | Unlimited                                                             |                           |                                            |    |     |     | Unlimited                       |                           |                                            |    |     |     |
| Security Contexts                      | No support                                                            |                           |                                            |    |     |     | 2                               | <i>Optional Licenses:</i> |                                            |    |     |     |
|                                        |                                                                       |                           |                                            |    |     |     |                                 | 5                         |                                            |    |     |     |
| VPN Sessions <sup>1</sup>              | 250 combined IPSec and WebVPN                                         |                           |                                            |    |     |     | 250 combined IPSec and WebVPN   |                           |                                            |    |     |     |
| Max. IPSec Sessions                    | 250                                                                   |                           |                                            |    |     |     | 250                             |                           |                                            |    |     |     |
| Max. WebVPN Sessions                   | 2                                                                     | <i>Optional Licenses:</i> |                                            |    |     |     | 2                               | <i>Optional Licenses:</i> |                                            |    |     |     |
|                                        |                                                                       | 10                        | 25                                         | 50 | 100 | 250 |                                 | 10                        | 25                                         | 50 | 100 | 250 |
| VPN Load Balancing                     | No support                                                            |                           |                                            |    |     |     | No support                      |                           |                                            |    |     |     |
| Failover                               | None                                                                  |                           |                                            |    |     |     | Active/Standby or Active/Active |                           |                                            |    |     |     |
| GTP/GPRS                               | No support                                                            |                           |                                            |    |     |     | No support                      |                           |                                            |    |     |     |
| Max. VLANs                             | 50                                                                    |                           |                                            |    |     |     | 100                             |                           |                                            |    |     |     |
| Concurrent Firewall Conns <sup>2</sup> | 50 K                                                                  |                           |                                            |    |     |     | 130 K                           |                           |                                            |    |     |     |
| Max. Physical Interfaces               | 3 at 10/100 plus the Management interface for management traffic only |                           |                                            |    |     |     | Unlimited                       |                           |                                            |    |     |     |
| Encryption                             | Base (DES)                                                            |                           | <i>Optional license: Strong (3DES/AES)</i> |    |     |     | Base (DES)                      |                           | <i>Optional license: Strong (3DES/AES)</i> |    |     |     |
| Min. RAM                               | 256 MB                                                                |                           |                                            |    |     |     | 256 MB                          |                           |                                            |    |     |     |

1. Although the maximum IPSec and WebVPN sessions add up to more than the maximum VPN sessions, the combined sessions should not exceed the VPN session limit. If you exceed the maximum VPN sessions, you can overload the security appliance, so be sure to size your network appropriately.
2. The concurrent firewall connections are based on a traffic mix of 80% TCP and 20% UDP, with 1 host and 1 dynamic translation for every 4 connections.

**Table A-3 ASA 5520 Adaptive Security Appliance License Features**

| ASA 5520                               | Base License                    |                    |                                     |    |     |     |           |     |
|----------------------------------------|---------------------------------|--------------------|-------------------------------------|----|-----|-----|-----------|-----|
| Users, concurrent                      | Unlimited                       |                    |                                     |    |     |     | Unlimited |     |
| Security Contexts                      | 2                               | Optional Licenses: |                                     |    |     |     |           |     |
|                                        |                                 | 5                  | 10                                  | 20 |     |     |           |     |
| VPN Sessions <sup>1</sup>              | 750 combined IPSec and WebVPN   |                    |                                     |    |     |     |           |     |
| Max. IPSec Sessions                    | 750                             |                    |                                     |    |     |     |           |     |
| Max. WebVPN Sessions                   | 2                               | Optional Licenses: |                                     |    |     |     |           |     |
|                                        |                                 | 10                 | 25                                  | 50 | 100 | 250 | 500       | 750 |
| VPN Load Balancing                     | Supported                       |                    |                                     |    |     |     |           |     |
| Failover                               | Active/Standby or Active/Active |                    |                                     |    |     |     |           |     |
| GTP/GPRS                               | None                            |                    | Optional license: Enabled           |    |     |     |           |     |
| Max. VLANs                             | 150                             |                    |                                     |    |     |     |           |     |
| Concurrent Firewall Conns <sup>2</sup> | 280 K                           |                    |                                     |    |     |     |           |     |
| Max. Physical Interfaces               | Unlimited                       |                    |                                     |    |     |     |           |     |
| Encryption                             | Base (DES)                      |                    | Optional license: Strong (3DES/AES) |    |     |     |           |     |
| Min. RAM                               | 512 MB                          |                    |                                     |    |     |     |           |     |

1. Although the maximum IPSec and WebVPN sessions add up to more than the maximum VPN sessions, the combined sessions should not exceed the VPN session limit. If you exceed the maximum VPN sessions, you can overload the security appliance, so be sure to size your network appropriately.
2. The concurrent firewall connections are based on a traffic mix of 80% TCP and 20% UDP, with 1 host and 1 dynamic translation for every 4 connections.

**Table A-4 ASA 5540 Adaptive Security Appliance License Features**

| ASA 5540                               | Base License                    |                    |                                     |    |     |     |           |     |      |      |
|----------------------------------------|---------------------------------|--------------------|-------------------------------------|----|-----|-----|-----------|-----|------|------|
| Users, concurrent                      | Unlimited                       |                    |                                     |    |     |     | Unlimited |     |      |      |
| Security Contexts                      | 2                               | Optional licenses: |                                     |    |     |     |           |     |      |      |
|                                        |                                 | 5                  | 10                                  | 20 | 50  |     |           |     |      |      |
| VPN Sessions <sup>1</sup>              | 5000 combined IPSec and WebVPN  |                    |                                     |    |     |     |           |     |      |      |
| Max. IPSec Sessions                    | 5000                            |                    |                                     |    |     |     |           |     |      |      |
| Max. WebVPN Sessions                   | 2                               | Optional Licenses: |                                     |    |     |     |           |     |      |      |
|                                        |                                 | 10                 | 25                                  | 50 | 100 | 250 | 500       | 750 | 1000 | 2500 |
| VPN Load Balancing                     | Supported                       |                    |                                     |    |     |     |           |     |      |      |
| Failover                               | Active/Standby or Active/Active |                    |                                     |    |     |     |           |     |      |      |
| GTP/GPRS                               | None                            |                    | Optional license: Enabled           |    |     |     |           |     |      |      |
| Max. VLANs                             | 200                             |                    |                                     |    |     |     |           |     |      |      |
| Concurrent Firewall Conns <sup>2</sup> | 400 K                           |                    |                                     |    |     |     |           |     |      |      |
| Max. Physical Interfaces               | Unlimited                       |                    |                                     |    |     |     |           |     |      |      |
| Encryption                             | Base (DES)                      |                    | Optional license: Strong (3DES/AES) |    |     |     |           |     |      |      |
| Min. RAM                               | 1 GB                            |                    |                                     |    |     |     |           |     |      |      |

1. Although the maximum IPSec and WebVPN sessions add up to more than the maximum VPN sessions, the combined sessions should not exceed the VPN session limit. If you exceed the maximum VPN sessions, you can overload the security appliance, so be sure to size your network appropriately.
2. The concurrent firewall connections are based on a traffic mix of 80% TCP and 20% UDP, with 1 host and 1 dynamic translation for every 4 connections.

**Table A-5 ASA 5550 Adaptive Security Appliance License Features**

| ASA 5550                               | Base License                    |                    |                                     |    |     |     |     |     |      |      |      |
|----------------------------------------|---------------------------------|--------------------|-------------------------------------|----|-----|-----|-----|-----|------|------|------|
| Users, concurrent                      | Unlimited                       |                    |                                     |    |     |     |     |     |      |      |      |
| Security Contexts                      | 2                               | Optional licenses: |                                     |    |     |     |     |     |      |      |      |
|                                        |                                 | 5                  | 10                                  | 20 | 50  |     |     |     |      |      |      |
| VPN Sessions <sup>1</sup>              | 5000 combined IPSec and WebVPN  |                    |                                     |    |     |     |     |     |      |      |      |
| Max. IPSec Sessions                    | 5000                            |                    |                                     |    |     |     |     |     |      |      |      |
| Max. WebVPN Sessions                   | 2                               | Optional Licenses: |                                     |    |     |     |     |     |      |      |      |
|                                        |                                 | 10                 | 25                                  | 50 | 100 | 250 | 500 | 750 | 1000 | 2500 | 5000 |
| VPN Load Balancing                     | Supported                       |                    |                                     |    |     |     |     |     |      |      |      |
| Failover                               | Active/Standby or Active/Active |                    |                                     |    |     |     |     |     |      |      |      |
| GTP/GPRS                               | None                            |                    | Optional license: Enabled           |    |     |     |     |     |      |      |      |
| Max. VLANs                             | 250                             |                    |                                     |    |     |     |     |     |      |      |      |
| Concurrent Firewall Conns <sup>2</sup> | 650 K                           |                    |                                     |    |     |     |     |     |      |      |      |
| Max. Physical Interfaces               | Unlimited                       |                    |                                     |    |     |     |     |     |      |      |      |
| Encryption                             | Base (DES)                      |                    | Optional license: Strong (3DES/AES) |    |     |     |     |     |      |      |      |
| Min. RAM                               | 4 GB                            |                    |                                     |    |     |     |     |     |      |      |      |

1. Although the maximum IPSec and WebVPN sessions add up to more than the maximum VPN sessions, the combined sessions should not exceed the VPN session limit. If you exceed the maximum VPN sessions, you can overload the security appliance, so be sure to size your network appropriately.
2. The concurrent firewall connections are based on a traffic mix of 80% TCP and 20% UDP, with 1 host and 1 dynamic translation for every 4 connections.

**Table A-6 PIX 515/515E Security Appliance License Features**

| PIX 515/515E                           | R (Restricted) |                              |                   | UR (Unrestricted)               |                              |                   | FO (Failover) <sup>1</sup> |                              |                   | FO-AA (Failover Active/Active) <sup>1</sup> |                              |                   |
|----------------------------------------|----------------|------------------------------|-------------------|---------------------------------|------------------------------|-------------------|----------------------------|------------------------------|-------------------|---------------------------------------------|------------------------------|-------------------|
| Users, concurrent                      | Unlimited      |                              |                   | Unlimited                       |                              |                   | Unlimited                  |                              |                   | Unlimited                                   |                              |                   |
| Security Contexts                      | No support     |                              |                   | 2                               | Optional license: 5          |                   | 2                          | Optional license: 5          |                   | 2                                           | Optional license: 5          |                   |
| IPSec Sessions                         | 2000           |                              |                   | 2000                            |                              |                   | 2000                       |                              |                   | 2000                                        |                              |                   |
| WebVPN Sessions                        | No support     |                              |                   | No support                      |                              |                   | No support                 |                              |                   | No support                                  |                              |                   |
| VPN Load Balancing                     | No support     |                              |                   | No support                      |                              |                   | No support                 |                              |                   | No support                                  |                              |                   |
| Failover                               | No support     |                              |                   | Active/Standby<br>Active/Active |                              |                   | Active/Standby             |                              |                   | Active/Standby<br>Active/Active             |                              |                   |
| GTP/GPRS                               | None           | Optional license:<br>Enabled |                   | None                            | Optional license:<br>Enabled |                   | None                       | Optional license:<br>Enabled |                   | None                                        | Optional license:<br>Enabled |                   |
| Max. VLANs                             | 10             |                              |                   | 25                              |                              |                   | 25                         |                              |                   | 25                                          |                              |                   |
| Concurrent Firewall Conns <sup>2</sup> | 48 K           |                              |                   | 130 K                           |                              |                   | 130 K                      |                              |                   | 130 K                                       |                              |                   |
| Max. Physical Interfaces               | 3              |                              |                   | 6                               |                              |                   | 6                          |                              |                   | 6                                           |                              |                   |
| Encryption                             | None           | Optional licenses:           |                   | None                            | Optional licenses:           |                   | None                       | Optional licenses:           |                   | None                                        | Optional licenses:           |                   |
|                                        |                | Base (DES)                   | Strong (3DES/AES) |                                 | Base (DES)                   | Strong (3DES/AES) |                            | Base (DES)                   | Strong (3DES/AES) |                                             | Base (DES)                   | Strong (3DES/AES) |
| Min. RAM                               | 64 MB          |                              |                   | 128 MB                          |                              |                   | 128 MB                     |                              |                   | 128 MB                                      |                              |                   |

1. This license can only be used in a failover pair with another unit with a UR license. Both units must be the same model.

2. The concurrent firewall connections are based on a traffic mix of 80% TCP and 20% UDP, with 1 host and 1 dynamic translation for every 4 connections.

**Table A-7 PIX 525 Security Appliance License Features**

| PIX 525                                | R (Restricted) |                              |                   | UR (Unrestricted)               |                              |                   |            | FO (Failover) <sup>1</sup> |                              |            |                   | FO-AA (Failover Active/Active) <sup>1</sup> |                              |    |   |
|----------------------------------------|----------------|------------------------------|-------------------|---------------------------------|------------------------------|-------------------|------------|----------------------------|------------------------------|------------|-------------------|---------------------------------------------|------------------------------|----|---|
| Users, concurrent                      | Unlimited      |                              |                   | Unlimited                       |                              |                   |            | Unlimited                  |                              |            |                   | Unlimited                                   |                              |    |   |
| Security Contexts                      | No support     |                              |                   | 2                               | Optional licenses:           |                   |            | 2                          | Optional licenses:           |            |                   | 2                                           | Optional licenses:           |    |   |
|                                        |                |                              |                   |                                 | 5                            | 10                | 20         |                            | 50                           | 5          | 10                |                                             | 20                           | 50 | 5 |
| IPSec Sessions                         | 2000           |                              |                   | 2000                            |                              |                   |            | 2000                       |                              |            |                   | 2000                                        |                              |    |   |
| WebVPN Sessions                        | No support     |                              |                   | No support                      |                              |                   |            | No support                 |                              |            |                   | No support                                  |                              |    |   |
| VPN Load Balancing                     | No support     |                              |                   | No support                      |                              |                   |            | No support                 |                              |            |                   | No support                                  |                              |    |   |
| Failover                               | No support     |                              |                   | Active/Standby<br>Active/Active |                              |                   |            | Active/Standby             |                              |            |                   | Active/Standby<br>Active/Active             |                              |    |   |
| GTP/GPRS                               | None           | Optional license:<br>Enabled |                   | None                            | Optional license:<br>Enabled |                   |            | None                       | Optional license:<br>Enabled |            |                   | None                                        | Optional license:<br>Enabled |    |   |
| Max. VLANs                             | 25             |                              |                   | 100                             |                              |                   |            | 100                        |                              |            |                   | 100                                         |                              |    |   |
| Concurrent Firewall Conns <sup>2</sup> | 140 K          |                              |                   | 280 K                           |                              |                   |            | 280 K                      |                              |            |                   | 280 K                                       |                              |    |   |
| Max. Physical Interfaces               | 6              |                              |                   | 10                              |                              |                   |            | 10                         |                              |            |                   | 10                                          |                              |    |   |
| Encryption                             | None           | Optional licenses:           |                   | None                            | Optional licenses:           |                   |            | None                       | Optional licenses:           |            |                   | None                                        | Optional licenses:           |    |   |
|                                        |                | Base (DES)                   | Strong (3DES/AES) |                                 | Base (DES)                   | Strong (3DES/AES) | Base (DES) |                            | Strong (3DES/AES)            | Base (DES) | Strong (3DES/AES) |                                             |                              |    |   |
| Min. RAM                               | 128 MB         |                              |                   | 256 MB                          |                              |                   |            | 256 MB                     |                              |            |                   | 256 MB                                      |                              |    |   |

1. This license can only be used in a failover pair with another unit with a UR license. Both units must be the same model.

2. The concurrent firewall connections are based on a traffic mix of 80% TCP and 20% UDP, with 1 host and 1 dynamic translation for every 4 connections.

**Table A-8 PIX 535 Security Appliance License Features**

| PIX 535                                | R (Restricted) |                              |                   | UR (Unrestricted)               |                              |                   |    | FO (Failover) <sup>1</sup> |                              |    |            | FO-AA (Failover Active/Active) <sup>1</sup> |                              |            |                   |
|----------------------------------------|----------------|------------------------------|-------------------|---------------------------------|------------------------------|-------------------|----|----------------------------|------------------------------|----|------------|---------------------------------------------|------------------------------|------------|-------------------|
| Users, concurrent                      | Unlimited      |                              |                   | Unlimited                       |                              |                   |    | Unlimited                  |                              |    |            | Unlimited                                   |                              |            |                   |
| Security Contexts                      | No support     |                              |                   | 2                               | Optional licenses:           |                   |    | 2                          | Optional licenses:           |    |            | 2                                           | Optional licenses:           |            |                   |
|                                        |                |                              |                   | 5                               | 10                           | 20                | 50 | 5                          | 10                           | 20 | 50         | 5                                           | 10                           | 20         | 50                |
| IPSec Sessions                         | 2000           |                              |                   | 2000                            |                              |                   |    | 2000                       |                              |    |            | 2000                                        |                              |            |                   |
| WebVPN Sessions                        | No support     |                              |                   | No support                      |                              |                   |    | No support                 |                              |    |            | No support                                  |                              |            |                   |
| VPN Load Balancing                     | No support     |                              |                   | No support                      |                              |                   |    | No support                 |                              |    |            | No support                                  |                              |            |                   |
| Failover                               | No support     |                              |                   | Active/Standby<br>Active/Active |                              |                   |    | Active/Standby             |                              |    |            | Active/Standby<br>Active/Active             |                              |            |                   |
| GTP/GPRS                               | None           | Optional license:<br>Enabled |                   | None                            | Optional license:<br>Enabled |                   |    | None                       | Optional license:<br>Enabled |    |            | None                                        | Optional license:<br>Enabled |            |                   |
| Max. VLANs                             | 50             |                              |                   | 150                             |                              |                   |    | 150                        |                              |    |            | 150                                         |                              |            |                   |
| Concurrent Firewall Conns <sup>2</sup> | 250 K          |                              |                   | 500 K                           |                              |                   |    | 500 K                      |                              |    |            | 500 K                                       |                              |            |                   |
| Max. Physical Interfaces               | 8              |                              |                   | 14                              |                              |                   |    | 14                         |                              |    |            | 14                                          |                              |            |                   |
| Encryption                             | None           | Optional licenses:           |                   | None                            | Optional licenses:           |                   |    | None                       | Optional licenses:           |    |            | None                                        | Optional licenses:           |            |                   |
|                                        |                | Base (DES)                   | Strong (3DES/AES) |                                 | Base (DES)                   | Strong (3DES/AES) |    | Base (DES)                 | Strong (3DES/AES)            |    | Base (DES) | Strong (3DES/AES)                           |                              | Base (DES) | Strong (3DES/AES) |
| Min. RAM                               | 512 MB         |                              |                   | 1024 MB                         |                              |                   |    | 1024 MB                    |                              |    |            | 1024 MB                                     |                              |            |                   |

1. This license can only be used in a failover pair with another unit with a UR license. Both units must be the same model.

2. The concurrent firewall connections are based on a traffic mix of 80% TCP and 20% UDP, with 1 host and 1 dynamic translation for every 4 connections.

## Security Services Module Support

Table A-9 shows the SSMs supported by each platform:

**Table A-9 SSM Support**

| Platform | SSM Models                                                      |
|----------|-----------------------------------------------------------------|
| ASA 5505 | No support                                                      |
| ASA 5510 | AIP SSM 10<br>AIP SSM 20<br>CSC SSM 10<br>CSC SSM 20<br>4GE SSM |

**Table A-9 SSM Support (continued)**

| Platform     | SSM Models                                                                                |
|--------------|-------------------------------------------------------------------------------------------|
| ASA 5520     | AIP SSM 10<br>AIP SSM 20<br>CSC SSM 10<br>CSC SSM 20<br>4GE SSM                           |
| ASA 5540     | AIP SSM 10<br>AIP SSM 20<br>CSC SSM 10 <sup>1</sup><br>CSC SSM 20 <sup>1</sup><br>4GE SSM |
| ASA 5550     | No support (4GE SSM is built-in and not user-removable)                                   |
| PIX 515/515E | No support                                                                                |
| PIX 525      | No support                                                                                |
| PIX 535      | No support                                                                                |

1. The CSC SSM licenses support up to 1000 users while the Cisco ASA 5540 Series appliance can support significantly more users. If you deploy CSC SSM with an ASA 5540 adaptive security appliance, be sure to configure the security appliance to send the CSC SSM only the traffic that should be scanned. For more information, see the [“Determining What Traffic to Scan” section on page 22-9](#) for more information.

## VPN Specifications

This section describes the VPN specifications for the security appliance. This section includes the following topics:

- [Cisco VPN Client Support, page A-11](#)
- [Cisco Secure Desktop Support, page A-11](#)
- [Site-to-Site VPN Compatibility, page A-11](#)
- [Cryptographic Standards, page A-12](#)



## Cisco VPN Client Support

The security appliance supports a wide variety of software and hardware-based Cisco VPN clients, as shown in [Table A-10](#).

**Table A-10** *Cisco VPN Client Support*

| Client Type                                        | Client Versions                                                          |
|----------------------------------------------------|--------------------------------------------------------------------------|
| SSL VPN clients                                    | Cisco SSL VPN client, Version 1.1 or higher                              |
| Software IPSec VPN clients                         | Cisco VPN client for Windows, Version 3.6 or higher                      |
|                                                    | Cisco VPN client for Linux, Version 3.6 or higher                        |
|                                                    | Cisco VPN client for Solaris, Version 3.6 or higher                      |
|                                                    | Cisco VPN client for Mac OS X, Version 3.6 or higher                     |
| Hardware IPSec VPN clients (Cisco Easy VPN remote) | Cisco VPN 3002 hardware client, Version 3.0 or higher                    |
|                                                    | Cisco IOS Software Easy VPN remote, Release 12.2(8)YJ                    |
|                                                    | Cisco PIX 500 series security appliance, Version 6.2 or higher           |
|                                                    | Cisco ASA 5500 series adaptive security appliance, Version 7.0 or higher |

## Cisco Secure Desktop Support

The security appliance supports CSD software Version 3.1.1.16.

## Site-to-Site VPN Compatibility

In addition to providing interoperability for many third-party VPN products, the security appliance interoperates with the Cisco VPN products for site-to-site VPN connectivity shown in [Table A-11](#).

**Table A-11** *Site-to-Site VPN Compatibility*

| Platforms                                          | Software Versions          |
|----------------------------------------------------|----------------------------|
| Cisco ASA 5500 series adaptive security appliances | Version 7.0(1) or higher   |
| Cisco IOS routers                                  | Release 12.1(6)T or higher |
| Cisco PIX 500 series security appliances           | Version 5.1(1) or higher   |
| Cisco VPN 3000 series concentrators                | Version 3.6(1) or higher   |

## Cryptographic Standards

The security appliance supports numerous cryptographic standards and related third-party products and services, including those shown in [Table A-12](#).

**Table A-12**      *Cryptographic Standards*

| Type                                                     | Description                                                                                                                                           |
|----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| Asymmetric (public key) encryption algorithms            | RSA public/private key pairs, 512 bits to 4096 bits<br>DSA public/private key pairs, 512 bits to 1024 bits                                            |
| Symmetric encryption algorithms                          | AES—128, 192, and 256 bits<br>DES—56 bits<br>3DES—168 bits<br>RC4—40, 56, 64, and 128 bits                                                            |
| Perfect forward secrecy (Diffie-Hellman key negotiation) | Group 1— 768 bits<br>Group 2—1024 bits<br>Group 5— 1536 bits<br>Group 7—163 bits (Elliptic Curve Diffie-Hellman)                                      |
| Hash algorithms                                          | MD5—128 bits<br>SHA-1—160 bits                                                                                                                        |
| X.509 certificate authorities                            | Cisco IOS software<br>Baltimore UniCERT<br>Entrust Authority<br>iPlanet/Netscape CMS<br>Microsoft Certificate Services<br>RSA Keon<br>VeriSign OnSite |
| X.509 certificate enrollment methods                     | SCEP<br>PKCS #7 and #10                                                                                                                               |



# APPENDIX **B**

## Sample Configurations

---

This appendix illustrates and describes a number of common ways to implement the security appliance, and includes the following topics:

- [Example 1: Multiple Mode Firewall With Outside Access, page B-1](#)
- [Example 2: Single Mode Firewall Using Same Security Level, page B-6](#)
- [Example 3: Shared Resources for Multiple Contexts, page B-8](#)
- [Example 4: Multiple Mode, Transparent Firewall with Outside Access, page B-12](#)
- [Example 5: WebVPN Configuration, page B-16](#)
- [Example 6: IPv6 Configuration, page B-18](#)
- [Example 7: Cable-Based Active/Standby Failover \(Routed Mode\), page B-20](#)
- [Example 8: LAN-Based Active/Standby Failover \(Routed Mode\), page B-21](#)
- [Example 9: LAN-Based Active/Active Failover \(Routed Mode\), page B-22](#)
- [Example 10: Cable-Based Active/Standby Failover \(Transparent Mode\), page B-26](#)
- [Example 11: LAN-Based Active/Standby Failover \(Transparent Mode\), page B-27](#)
- [Example 12: LAN-Based Active/Active Failover \(Transparent Mode\), page B-28](#)
- [Example 13: Dual ISP Support Using Static Route Tracking, page B-31](#)
- [Example 14: ASA 5505 Base License, page B-33](#)
- [Example 15: ASA 5505 Security Plus License with Failover and Dual-ISP Backup, page B-35](#)
- [Example 16: Network Traffic Diversion, page B-37](#)

### Example 1: Multiple Mode Firewall With Outside Access

This configuration creates three security contexts plus the admin context, each with an inside and an outside interface. The Customer C context includes a DMZ interface where a Websense server for HTTP filtering resides on the service provider premises (see [Figure B-1](#)).


Inside hosts can access the Internet through the outside using dynamic NAT or PAT, but no outside hosts can access the inside.

The Customer A context has a second network behind an inside router.

The admin context allows SSH sessions to the security appliance from one host.

Although inside IP addresses can be the same across contexts when the interfaces are unique, keeping them unique is easier to manage.

**Figure B-1**      **Example 1**



See the following sections for the configurations for this scenario:

- [Example 1: System Configuration, page B-2](#)
- [Example 1: Admin Context Configuration, page B-4](#)
- [Example 1: Customer A Context Configuration, page B-4](#)
- [Example 1: Customer B Context Configuration, page B-4](#)
- [Example 1: Customer C Context Configuration, page B-5](#)

## Example 1: System Configuration

You must first enable multiple context mode using the **mode multiple** command. The mode is not stored in the configuration file, even though it endures reboots. Enter the **show mode** command to view the current mode.

```
hostname Farscape
password passw0rd
enable password chr1cht0n
mac-address auto
asdm image disk0:/asdm.bin
boot system disk0:/image.bin
admin-context admin
interface gigabitethernet 0/0
 shutdown
interface gigabitethernet 0/0.3
 vlan 3
 no shutdown
interface gigabitethernet 0/1
 no shutdown
interface gigabitethernet 0/1.4
 vlan 4
 no shutdown
interface gigabitethernet 0/1.5
 vlan 5
 no shutdown
interface gigabitethernet 0/1.6
 vlan 6
 no shutdown
interface gigabitethernet 0/1.7
 vlan 7
 no shutdown
interface gigabitethernet 0/1.8
 vlan 8
 no shutdown
class gold
 limit-resource rate conns 2000
 limit-resource conns 20000
class silver
 limit-resource rate conns 1000
 limit-resource conns 10000
class bronze
 limit-resource rate conns 500
 limit-resource conns 5000
context admin
 allocate-interface gigabitethernet 0/0.3
 allocate-interface gigabitethernet 0/1.4
 config-url disk0://admin.cfg
 member default
context customerA
 description This is the context for customer A
 allocate-interface gigabitethernet 0/0.3
 allocate-interface gigabitethernet 0/1.5
 config-url disk0://contexta.cfg
 member gold
context customerB
 description This is the context for customer B
 allocate-interface gigabitethernet 0/0.3
 allocate-interface gigabitethernet 0/1.6
 config-url disk0://contextb.cfg
 member silver
context customerC
 description This is the context for customer C
 allocate-interface gigabitethernet 0/0.3
 allocate-interface gigabitethernet 0/1.7-gigabitethernet 0/1.8
 config-url disk0://contextc.cfg
 member bronze
```

## Example 1: Admin Context Configuration

The host at 10.1.1.75 can access the context using SSH, which requires a key to be generated using the **crypto key generate** command.

```
hostname Admin
domain isp
interface gigabitethernet 0/0.3
 nameif outside
 security-level 0
 ip address 209.165.201.2 255.255.255.224
 no shutdown
interface gigabitethernet 0/1.4
 nameif inside
 security-level 100
 ip address 10.1.1.1 255.255.255.0
 no shutdown
passwd secret1969
enable password h1and10
route outside 0 0 209.165.201.1 1
ssh 10.1.1.75 255.255.255.255 inside
nat (inside) 1 10.1.1.0 255.255.255.0
! This context uses dynamic NAT for inside users that access the outside
global (outside) 1 209.165.201.10-209.165.201.29
! The host at 10.1.1.75 has access to the Websense server in Customer C, so
! it needs a static translation for use in Customer C's access list
static (inside,outside) 209.165.201.30 10.1.1.75 netmask 255.255.255.255
```

## Example 1: Customer A Context Configuration

```
interface gigabitethernet 0/0.3
 nameif outside
 security-level 0
 ip address 209.165.201.3 255.255.255.224
 no shutdown
interface gigabitethernet 0/1.5
 nameif inside
 security-level 100
 ip address 10.1.2.1 255.255.255.0
 no shutdown
passwd hell0!
enable password enter55
route outside 0 0 209.165.201.1 1
! The Customer A context has a second network behind an inside router that requires a
! static route. All other traffic is handled by the default route pointing to the router.
route inside 192.168.1.0 255.255.255.0 10.1.2.2 1
nat (inside) 1 10.1.2.0 255.255.255.0
! This context uses dynamic PAT for inside users that access that outside. The outside
! interface address is used for the PAT address
global (outside) 1 interface
```

## Example 1: Customer B Context Configuration

```
interface gigabitethernet 0/0.3
 nameif outside
 security-level 0
 ip address 209.165.201.4 255.255.255.224
```

```

no shutdown
interface gigabitethernet 0/1.6
 nameif inside
 security-level 100
 ip address 10.1.3.1 255.255.255.0
 no shutdown
passwd tenac1ous
enable password defen$e
route outside 0 0 209.165.201.1 1
nat (inside) 1 10.1.3.0 255.255.255.0
! This context uses dynamic PAT for inside users that access the outside
global (outside) 1 209.165.201.9 netmask 255.255.255.255
access-list INTERNET remark Inside users only access HTTP and HTTPS servers on the outside
access-list INTERNET extended permit tcp any any eq http
access-list INTERNET extended permit tcp any any eq https
access-group INTERNET in interface inside

```

## Example 1: Customer C Context Configuration

```

interface gigabitethernet 0/0.3
 nameif outside
 security-level 0
 ip address 209.165.201.5 255.255.255.224
 no shutdown
interface gigabitethernet 0/1.7
 nameif inside
 security-level 100
 ip address 10.1.4.1 255.255.255.0
 no shutdown
interface gigabitethernet 0/1.8
 nameif dmz
 security-level 50
 ip address 192.168.2.1 255.255.255.0
 no shutdown
passwd fl0wer
enable password treeh0u$e
route outside 0 0 209.165.201.1 1
url-server (dmz) vendor websense host 192.168.2.2 url-block block 50
url-cache dst 128
filter url http 10.1.4.0 255.255.255.0 0 0
! When inside users access an HTTP server, the security appliance consults with a
! Websense server to determine if the traffic is allowed
nat (inside) 1 10.1.4.0 255.255.255.0
! This context uses dynamic NAT for inside users that access the outside
global (outside) 1 209.165.201.9 netmask 255.255.255.255
! A host on the admin context requires access to the Websense server for management using
! pcAnywhere, so the Websense server uses a static translation for its private address
static (dmz,outside) 209.165.201.6 192.168.2.2 netmask 255.255.255.255
access-list MANAGE remark Allows the management host to use pcAnywhere on the Websense
server
access-list MANAGE extended permit tcp host 209.165.201.30 host 209.165.201.6 eq
pcanywhere-data
access-list MANAGE extended permit udp host 209.165.201.30 host 209.165.201.6 eq
pcanywhere-status
access-group MANAGE in interface outside


```

## Example 2: Single Mode Firewall Using Same Security Level

This configuration creates three internal interfaces. Two of the interfaces connect to departments that are on the same security level, which allows all hosts to communicate without using access lists. The DMZ interface hosts a Syslog server. The management host on the outside needs access to the Syslog server and the security appliance. To connect to the security appliance, the host uses a VPN connection. The security appliance uses RIP on the inside interfaces to learn routes. The security appliance does not advertise routes with RIP; the upstream router needs to use static routes for security appliance traffic (see [Figure B-2](#)).

The Department networks are allowed to access the Internet, and use PAT.

**Figure B-2** Example 2



```
passwd g00fba11
enable password genlu$
hostname Buster
asdm image disk0:/asdm.bin
boot system disk0:/image.bin
interface gigabitethernet 0/0
 nameif outside
 security-level 0
```



```

ip address 209.165.201.3 255.255.255.224
no shutdown
interface gigabitethernet 0/1
 nameif dept2
 security-level 100
 ip address 10.1.2.1 255.255.255.0
 mac-address 000C.F142.4CDE standby 000C.F142.4CDF
 no shutdown
 rip authentication mode md5
 rip authentication key scorpius key_id 1
interface gigabitethernet 0/2
 nameif dept1
 security-level 100
 ip address 10.1.1.1 255.255.255.0
 no shutdown
interface gigabitethernet 0/3
 nameif dmz
 security-level 50
 ip address 192.168.2.1 255.255.255.0
 no shutdown
same-security-traffic permit inter-interface
route outside 0 0 209.165.201.1 1
nat (dept1) 1 10.1.1.0 255.255.255.0
nat (dept2) 1 10.1.2.0 255.255.255.0
! The dept1 and dept2 networks use PAT when accessing the outside
global (outside) 1 209.165.201.9 netmask 255.255.255.255
! Because we perform dynamic NAT on these addresses for outside access, we need to perform
! NAT on them for all other interface access. This identity static statement just
! translates the local address to the same address.
static (dept1,dept2) 10.1.1.0 10.1.1.0 netmask 255.255.255.0
static (dept2,dept1) 10.1.2.0 10.1.2.0 netmask 255.255.255.0
! The syslog server uses a static translation so the outside management host can access
! the server
static (dmz,outside) 209.165.201.5 192.168.2.2 netmask 255.255.255.255
access-list MANAGE remark Allows the management host to access the syslog server
access-list MANAGE extended permit tcp host 209.165.200.225 host 209.165.201.5 eq telnet
access-group MANAGE in interface outside
! Advertises the security appliance IP address as the default gateway for the downstream
! router. The security appliance does not advertise a default route to the upstream
! router. Listens for RIP updates from the downstream router. The security appliance does
! not listen for RIP updates from the upstream router because a default route to the
! upstream router is all that is required.
router rip
 network 10.0.0.0
 default information originate
 version 2
! The client uses a pre-shared key to connect to the security appliance over IPsec. The
! key is the password in the username command following.
isakmp policy 1 authentication pre-share
isakmp policy 1 encryption 3des
isakmp policy 1 group 2
isakmp policy 1 hash sha
isakmp enable outside
crypto ipsec transform-set vpn_client esp-3des esp-sha-hmac
username admin password passw0rd
crypto ipsec transform-set vpn esp-3des esp-sha-hmac
crypto dynamic-map vpn_client 1 set transform-set vpn
crypto map telnet_tunnel 1 ipsec-isakmp dynamic vpn_client
crypto map telnet_tunnel interface outside
ip local pool client_pool 10.1.1.2
access-list VPN_SPLIT extended permit ip host 209.165.201.3 host 10.1.1.2
telnet 10.1.1.2 255.255.255.255 outside
telnet timeout 30
logging trap 5

```

### Example 3: Shared Resources for Multiple Contexts


```
! System messages are sent to the syslog server on the DMZ network
logging host dmz 192.168.2.2
logging enable
```

## Example 3: Shared Resources for Multiple Contexts

This configuration includes multiple contexts for multiple departments within a company. Each department has its own security context so that each department can have its own security policy. However, the syslog, mail, and AAA servers are shared across all departments. These servers are placed on a shared interface (see [Figure B-3](#)).

Department 1 has a web server that outside users who are authenticated by the AAA server can access.

**Figure B-3** Example 3



See the following sections for the configurations for this scenario:

- [Example 3: System Configuration, page B-9](#)
- [Example 3: Admin Context Configuration, page B-9](#)

- [Example 3: Department 1 Context Configuration, page B-10](#)
- [Example 3: Department 2 Context Configuration, page B-11](#)

## Example 3: System Configuration

You must first enable multiple context mode using the **mode multiple** command. The mode is not stored in the configuration file, even though it endures reboots. Enter the **show mode** command to view the current mode.

```
hostname Ubik
password pkd55
enable password deckard69
asdm image disk0:/asdm.bin
boot system disk0:/image.bin
mac-address auto
admin-context admin
interface gigabitethernet 0/0
 no shutdown
interface gigabitethernet 0/0.200
 vlan 200
 no shutdown
interface gigabitethernet 0/1
 shutdown
interface gigabitethernet 0/1.201
 vlan 201
 no shutdown
interface gigabitethernet 0/1.202
 vlan 202
 no shutdown
interface gigabitethernet 0/1.300
 vlan 300
 no shutdown
context admin
 allocate-interface gigabitethernet 0/0.200
 allocate-interface gigabitethernet 0/1.201
 allocate-interface gigabitethernet 0/1.300
 config-url disk0://admin.cfg
context department1
 allocate-interface gigabitethernet 0/0.200
 allocate-interface gigabitethernet 0/1.202
 allocate-interface gigabitethernet 0/1.300
 config-url ftp://admin:passw0rd@10.1.0.16/dept1.cfg
context department2
 allocate-interface gigabitethernet 0/0.200
 allocate-interface gigabitethernet 0/1.203
 allocate-interface gigabitethernet 0/1.300
 config-url ftp://admin:passw0rd@10.1.0.16/dept2.cfg
```

## Example 3: Admin Context Configuration

```
hostname Admin
interface gigabitethernet 0/0.200
 nameif outside
 security-level 0
 ip address 209.165.201.3 255.255.255.224
 no shutdown
interface gigabitethernet 0/0.201
 nameif inside
```

### Example 3: Shared Resources for Multiple Contexts

```

security-level 100
ip address 10.1.0.1 255.255.255.0
no shutdown
interface gigabitethernet 0/0.300
nameif shared
security-level 50
ip address 10.1.1.1 255.255.255.0
no shutdown
passwd v00d00
enable password d011
route outside 0 0 209.165.201.2 1
nat (inside) 1 10.1.0.0 255.255.255.0
! This context uses PAT for inside users that access the outside
global (outside) 1 209.165.201.6 netmask 255.255.255.255
! This context uses PAT for inside users that access the shared network
global (shared) 1 10.1.1.30
! Because this host can access the web server in the Department 1 context, it requires a
! static translation
static (inside,outside) 209.165.201.7 10.1.0.15 netmask 255.255.255.255
! Because this host has management access to the servers on the Shared interface, it
! requires a static translation to be used in an access list
static (inside,shared) 10.1.1.78 10.1.0.15 netmask 255.255.255.255
access-list SHARED remark -Allows only mail traffic from inside to exit shared interface
access-list SHARED remark -but allows the admin host to access any server.
access-list SHARED extended permit ip host 10.1.1.78 any
access-list SHARED extended permit tcp host 10.1.1.30 host 10.1.1.7 eq smtp
! Note that the translated addresses are used.
access-group SHARED out interface shared
! Allows 10.1.0.15 to access the admin context using Telnet. From the admin context, you
! can access all other contexts.
telnet 10.1.0.15 255.255.255.255 inside
aaa-server AAA-SERVER protocol tacacs+
aaa-server AAA-SERVER (shared) host 10.1.1.6
key TheUauthKey
server-port 16
! The host at 10.1.0.15 must authenticate with the AAA server to log in
aaa authentication telnet console AAA-SERVER
aaa authorization command AAA-SERVER LOCAL
aaa accounting command AAA-SERVER
logging trap 6
! System messages are sent to the syslog server on the Shared network
logging host shared 10.1.1.8
logging enable

```

### Example 3: Department 1 Context Configuration

```

interface gigabitethernet 0/0.200
nameif outside
security-level 0
ip address 209.165.201.4 255.255.255.224
no shutdown
interface gigabitethernet 0/0.202
nameif inside
security-level 100
ip address 10.1.2.1 255.255.255.0
no shutdown
interface gigabitethernet 0/0.300
nameif shared
security-level 50
ip address 10.1.1.2 255.255.255.0
no shutdown

```

```

passwd cugel
enable password rhialto
nat (inside) 1 10.1.2.0 255.255.255.0
! The inside network uses PAT when accessing the outside
global (outside) 1 209.165.201.8 netmask 255.255.255.255
! The inside network uses dynamic NAT when accessing the shared network
global (shared) 1 10.1.1.31-10.1.1.37
! The web server can be accessed from outside and requires a static translation
static (inside,outside) 209.165.201.9 10.1.2.3 netmask 255.255.255.255
access-list WEBSERVER remark -Allows the management host (its translated address) on the
access-list WEBSERVER remark -admin context to access the web server for management
access-list WEBSERVER remark -it can use any IP protocol
access-list WEBSERVER extended permit ip host 209.165.201.7 host 209.165.201.9
access-list WEBSERVER remark -Allows any outside address to access the web server
access-list WEBSERVER extended permit tcp any eq http host 209.165.201.9 eq http
access-group WEBSERVER in interface outside
access-list MAIL remark -Allows only mail traffic from inside to exit out the shared int
! Note that the translated addresses are used.
access-list MAIL extended permit tcp host 10.1.1.31 eq smtp host 10.1.1.7 eq smtp
access-list MAIL extended permit tcp host 10.1.1.32 eq smtp host 10.1.1.7 eq smtp
access-list MAIL extended permit tcp host 10.1.1.33 eq smtp host 10.1.1.7 eq smtp
access-list MAIL extended permit tcp host 10.1.1.34 eq smtp host 10.1.1.7 eq smtp
access-list MAIL extended permit tcp host 10.1.1.35 eq smtp host 10.1.1.7 eq smtp
access-list MAIL extended permit tcp host 10.1.1.36 eq smtp host 10.1.1.7 eq smtp
access-list MAIL extended permit tcp host 10.1.1.37 eq smtp host 10.1.1.7 eq smtp
access-group MAIL out interface shared
aaa-server AAA-SERVER protocol tacacs+
aaa-server AAA-SERVER (shared) host 10.1.1.6
 key TheUauthKey
 server-port 16
! All traffic matching the WEBSERVER access list must authenticate with the AAA server
aaa authentication match WEBSERVER outside AAA-SERVER
logging trap 4
! System messages are sent to the syslog server on the Shared network
logging host shared 10.1.1.8
logging enable

```

### Example 3: Department 2 Context Configuration

```

interface gigabitethernet 0/0.200
 nameif outside
 security-level 0
 ip address 209.165.201.5 255.255.255.224
 no shutdown
interface gigabitethernet 0/0.203
 nameif inside
 security-level 100
 ip address 10.1.3.1 255.255.255.0
 no shutdown
interface gigabitethernet 0/0.300
 nameif shared
 security-level 50
 ip address 10.1.1.3 255.255.255.0
 no shutdown
passwd mazlrlan
enable password ly0ne$$e
route outside 0 0 209.165.201.2 1
nat (inside) 1 10.1.3.0 255.255.255.0
! The inside network uses PAT when accessing the outside
global (outside) 1 209.165.201.10 netmask 255.255.255.255
! The inside network uses PAT when accessing the shared network

```

```
global (shared) 1 10.1.1.38
access-list MAIL remark -Allows only mail traffic from inside to exit out the shared int
access-list MAIL extended permit tcp host 10.1.1.38 host 10.1.1.7 eq smtp
! Note that the translated PAT address is used.
access-group MAIL out interface shared
logging trap 3
! System messages are sent to the syslog server on the Shared network
logging host shared 10.1.1.8
logging enable
```

## Example 4: Multiple Mode, Transparent Firewall with Outside Access

This configuration creates three security contexts plus the admin context. Each context allows OSPF traffic to pass between the inside and outside routers (see [Figure B-4](#)).


Inside hosts can access the Internet through the outside, but no outside hosts can access the inside.

An out-of-band management host is connected to the Management 0/0 interface.

The admin context allows SSH sessions to the security appliance from one host.

Although inside IP addresses can be the same across contexts, keeping them unique is easier to manage.

Figure B-4 Example 4



See the following sections for the configurations for this scenario:

- [Example 4: System Configuration, page B-13](#)
- [Example 4: Admin Context Configuration, page B-14](#)
- [Example 4: Customer A Context Configuration, page B-15](#)
- [Example 4: Customer B Context Configuration, page B-15](#)
- [Example 4: Customer C Context Configuration, page B-16](#)

## Example 4: System Configuration

You must first enable multiple context mode using the **mode multiple** command. The mode is not stored in the configuration file, even though it endures reboots. Enter the **show mode** command to view the current mode.

```
firewall transparent
hostname Farscape
password passw0rd
enable password chr1cht0n
asdm image disk0:/asdm.bin
boot system disk0:/image.bin
admin-context admin
interface gigabitethernet 0/0
```

#### Example 4: Multiple Mode, Transparent Firewall with Outside Access

```

no shutdown
interface gigabitethernet 0/0.150
 vlan 150
 no shutdown
interface gigabitethernet 0/0.151
 vlan 151
 no shutdown
interface gigabitethernet 0/0.152
 vlan 152
 no shutdown
interface gigabitethernet 0/0.153
 vlan 153
 no shutdown
interface gigabitethernet 0/1
 shutdown
interface gigabitethernet 0/1.4
 vlan 4
 no shutdown
interface gigabitethernet 0/1.5
 vlan 5
 no shutdown
interface gigabitethernet 0/1.6
 vlan 6
 no shutdown
interface gigabitethernet 0/1.7
 vlan 7
 no shutdown
interface management 0/0
 no shutdown
context admin
 allocate-interface gigabitethernet 0/0.150
 allocate-interface gigabitethernet 0/1.4
 allocate-interface management 0/0
 config-url disk0://admin.cfg
context customerA
 description This is the context for customer A
 allocate-interface gigabitethernet 0/0.151
 allocate-interface gigabitethernet 0/1.5
 config-url disk0://contexta.cfg
context customerB
 description This is the context for customer B
 allocate-interface gigabitethernet 0/0.152
 allocate-interface gigabitethernet 0/1.6
 config-url disk0://contextb.cfg
context customerC
 description This is the context for customer C
 allocate-interface gigabitethernet 0/0.153
 allocate-interface gigabitethernet 0/1.7
 config-url disk0://contextc.cfg

```

## Example 4: Admin Context Configuration

The host at 10.1.1.75 can access the context using SSH, which requires a key pair to be generated using the **crypto key generate** command.

```

hostname Admin
domain isp
interface gigabitethernet 0/0.150
 nameif outside
 security-level 0
 no shutdown

```



```
interface gigabitethernet 0/1.4
 nameif inside
 security-level 100
 no shutdown
interface management 0/0
 nameif manage
 security-level 50
 ip address 10.2.1.1 255.255.255.0
 no shutdown
passwd secret1969
enable password hland10
ip address 10.1.1.1 255.255.255.0
route outside 0 0 10.1.1.2 1
ssh 10.1.1.75 255.255.255.255 inside
access-list OSPF remark -Allows OSPF
access-list OSPF extended permit 89 any any
access-group OSPF in interface outside
```

## Example 4: Customer A Context Configuration

```
interface gigabitethernet 0/0.151
 nameif outside
 security-level 0
 no shutdown
interface gigabitethernet 0/1.5
 nameif inside
 security-level 100
 no shutdown
passwd hell0!
enable password enter55
ip address 10.1.2.1 255.255.255.0
route outside 0 0 10.1.2.2 1
access-list OSPF remark -Allows OSPF
access-list OSPF extended permit 89 any any
access-group OSPF in interface outside
```

## Example 4: Customer B Context Configuration

```
interface gigabitethernet 0/0.152
 nameif outside
 security-level 0
 no shutdown
interface gigabitethernet 0/1.6
 nameif inside
 security-level 100
 no shutdown
passwd tenac10us
enable password defen$e
ip address 10.1.3.1 255.255.255.0
route outside 0 0 10.1.3.2 1
access-list OSPF remark -Allows OSPF
access-list OSPF extended permit 89 any any
access-group OSPF in interface outside
```

## Example 4: Customer C Context Configuration

```
interface gigabitethernet 0/0.153
 nameif outside
 security-level 0
 no shutdown
interface gigabitethernet 0/1.7
 nameif inside
 security-level 100
 no shutdown
passwd f10wer
enable password treeh0u$e
ip address 10.1.4.1 255.255.255.0
route outside 0 0 10.1.4.2 1
access-list OSPF remark -Allows OSPF
access-list OSPF extended permit 89 any any
access-group OSPF in interface outside
```

## Example 5: WebVPN Configuration

This configuration shows the commands needed to create WebVPN connections to the security appliance.

WebVPN lets users establish a secure, remote-access VPN tunnel to the security appliance using a web browser. There is no need for either a software or hardware client. WebVPN provides easy access to a broad range of web resources and web-enabled applications from almost any computer that can reach HTTP(S) Internet sites. WebVPN uses Secure Socket Layer Protocol and its successor, Transport Layer Security (SSL/TLS1) to provide a secure connection between remote users and specific, supported internal resources that you configure at a central site. The security appliance recognizes connections that need to be proxied, and the HTTP server interacts with the authentication subsystem to authenticate users.

---

### Step 1 Configure the security appliance for WebVPN.

```
webvpn
! WebVPN sessions are allowed on the outside and dmz1 interfaces, ASDM is not allowed.
enable outside
enable dmz161
title-color green
secondary-color 200,160,0
text-color black
default-idle-timeout 3600
! The NetBios Name server used for CIFS resolution.
nbns-server 172.31.122.10 master timeout 2 retry 2
accounting-server-group RadiusACS1
! WebVPN sessions are authenticated to a RADIUS aaa server.
authentication-server-group RadiusACS2
```

### Step 2 You must enable WebVPN access lists to be enforced on a group-policy or user policy. The access lists are defined with the **filter value** and **functions** commands in the group or user configuration.

```
access-list maia2 remark -deny access to url and send a syslog every 300 seconds
access-list maia2 remark -containing the hit-count (how many times the url was accessed)
access-list maia2 webtype deny url https://sales.example.com log informational interval 300
access-list maia2 remark -Permits access to the URL.
access-list maia2 webtype permit url http://employee-connection.example.com
```

```
access-list maia2 remark -Permits access to the site using ssh.
access-list maia2 remark -To be enforced via Port-Forwarding application.
access-list maia2 webtype permit tcp asa-35.example.com 255.255.255.255 eq ssh
access-list maia2 remark -Denies access to the application on port 1533.
access-list maia2 webtype deny tcp im.example.com 255.255.255.255 eq 1533
access-list maia2 remark -Permits access to files on this file share via
access-list maia2 remark -WebVPN Common Internet File System (CIFS).
access-list maia2 webtype permit url cifs://server-bos/people/mkting log informational
3600
```

- Step 3** You can configure a list of pre-configured URLs presented on the WebVPN user's home page after login, which are defined per user or per group.

```
url-list HomeURL "Sales" https://sales.example.com
url-list HomeURL "VPN3000-1" http://vpn3k-1.example.com
url-list HomeURL "OWA-2000" http://10.160.105.2/exchange
url-list HomeURL "Exchange5.5" http://10.86.195.113/exchange
url-list HomeURL "Employee Benefits" http://benefits.example.com
url-list HomeURL "Calendar" http://http://eng.example.com/cal.html
```

- Step 4** Configure a list of non-web TCP applications that will be port-forwarded over WebVPN and enforced per user or per group-policy. These are defined globally but can be enforced per user or per group-policy.

```
port-forward Apps1 4001 10.148.1.81 telnet term-srvr
port-forward Apps1 4008 router1-example.com ssh
port-forward Apps1 10143 flask.example.com imap4
port-forward Apps1 10110 flask.example.com pop3
port-forward Apps1 10025 flask.example.com smtp
port-forward Apps1 11533 sametime-im.example.com 1533
port-forward Apps1 10022 secure-term.example.com ssh
port-forward Apps1 21666 tuscan.example.com 1666 perforce-f1
port-forward Apps1 1030 sales.example.com https
```

- Step 5** Configure the policy attributes enforced for users of the SSLVPNusers group-policy.

```
group-policy SSLVPNusers internal
group-policy SSLVPNusers attributes
 banner value Welcome to Web Services !!!
 vpn-idle-timeout 2
 vpn-tunnel-protocol IPSec webvpn
 webvpn
 functions url-entry file-access file-entry file-browsing port-forward filter
 url-list value HomeURL
 port-forward value Apps1
```

- Step 6** Next, configure the interface(s) where ASDM and WebVPN HTTPS sessions will terminate. Note that The security appliance can support both WebVPN and an ASDM administrative session simultaneously on the same interface. To do so, you must assign different port numbers to these functions.

```
! Enables the HTTP server to allow ASDM and WebVPN HTTPS sessions.
http server enable
! Allows ASDM session(s) from host 10.20.30.47 on the inside interface ; WebVPN sessions
! are not allowed on this interface.
http 10.10.10.45 255.255.255.255 inside
! Allows WebVPN sessions on outside interface using HTTP to be re-directed to HTTPS.
! ASDM session is not allowed on this interface.
http redirect outside 80
! Allows WebVPN sessions on dmz1 interface using HTTP to be re-directed to HTTPS.
http redirect dmz1 80
```

- Step 7** Next, allow HTTPS ASDM and WebVPN sessions to terminate on the security appliance using the 3DES-sha1 cipher. Requires that a proper 3DES activation-key be previously installed.

```
ssl encryption 3des-sha1
```

```
ssl trust-point CA-MS inside
```


**Step 8** Finally, configure the email proxy settings.

```
imap4s
 enable outside
 enable inside
 enable dmz161
 default-group-policy DfltGrpPolicy
pop3s
 enable outside
 enable inside
 enable dmz161
 default-group-policy DfltGrpPolicy
smtps
 enable outside
 enable inside
 enable dmz161
 default-group-policy DfltGrpPolicy
```

## Example 6: IPv6 Configuration

This sample configuration shows several features of IPv6 support on the security appliance:

- Each interface is configured with both IPv6 and IPv4 addresses.
- The IPv6 default route is set with the **ipv6 route** command.
- An IPv6 access list is applied to the outside interface.
- The enforcement of Modified-EUI64 format interface identifiers in the IPv6 addresses of hosts on the inside interface.
- The outside interface suppresses router advertisement messages.
- An IPv6 static route.

**Figure B-5 IPv6 Dual Stack Configuration**

153155

```

enable password myenablepassword
passwd mypassword
hostname coyupix
asdm image flash:/asdm.bin
boot system flash:/image.bin
interface Ethernet0
 nameif outside
 security-level 0
 ip address 10.142.10.100 255.255.255.0
 ipv6 address 2001:400:3:1::100/64
 ipv6 nd suppress-ra
 ospf mtu-ignore auto
 no shutdown
interface Ethernet1
 nameif inside
 security-level 100
 ip address 10.140.10.100 255.255.255.0
 ipv6 address 2001:400:1:1::100/64
 ospf mtu-ignore auto
 no shutdown
access-list allow extended permit icmp any any
ssh 10.140.10.75 255.255.255.255 inside
logging enable
logging buffered debugging
ipv6 enforce-eui64 inside
ipv6 route outside 2001:400:6:1::/64 2001:400:3:1::1
ipv6 route outside ::/0 2001:400:3:1::1
ipv6 access-list outacl permit icmp6 2001:400:2:1::/64 2001:400:1:1::/64
ipv6 access-list outacl permit tcp 2001:400:2:1::/64 2001:400:1:1::/64 eq telnet
ipv6 access-list outacl permit tcp 2001:400:2:1::/64 2001:400:1:1::/64 eq ftp
ipv6 access-list outacl permit tcp 2001:400:2:1::/64 2001:400:1:1::/64 eq www
access-group allow in interface outside

```


## Example 7: Cable-Based Active/Standby Failover (Routed Mode)

```
access-group outacl in interface outside
route outside 0.0.0.0 0.0.0.0 16.142.10.1 1
```

# Example 7: Cable-Based Active/Standby Failover (Routed Mode)

Figure B-6 shows the network diagram for a failover configuration using a serial Failover cable. This configuration is only available on the PIX security appliance.

**Figure B-6** Cable-Based Failover Configuration



The following are the typical commands in a cable-based failover configuration.

```
enable password myenablepassword
passwd mypassword
hostname pixfirewall
asdm image flash:/asdm.bin
boot system flash:/image.bin
interface Ethernet0
 nameif outside
 security-level 0
 speed 100
 duplex full
 ip address 209.165.201.1 255.255.255.224 standby 209.165.201.2
 no shutdown
interface Ethernet1
 nameif inside
 security-level 100
 speed 100
 duplex full
 ip address 192.168.2.1 255.255.255.0 standby 192.168.2.2
 no shutdown
```

```


interface Ethernet3
 description STATE Failover Interface
 telnet 192.168.2.45 255.255.255.255 inside
 access-list acl_in permit tcp any host 209.165.201.5 eq 80
 access-group acl_in in interface outside
 failover
 failover link state Ethernet3
 failover interface ip state 192.168.253.1 255.255.255.252 standby 192.168.253.2
 global (outside) 1 209.165.201.3 netmask 255.255.255.224
 nat (inside) 1 0.0.0.0 0.0.0.0
 static (inside,outside) 209.165.201.5 192.168.2.5 netmask 255.255.255.255 0 0
 route outside 0.0.0.0 0.0.0.0 209.165.201.4 1

```

## Example 8: LAN-Based Active/Standby Failover (Routed Mode)

Figure B-7 shows the network diagram for a failover configuration using an Ethernet failover link. The units are configured to detect unit failures and to fail over in under a second (see the **failover polltime unit** command in the primary unit configuration).

**Figure B-7** LAN-Based Failover Configuration



See the following sections for the configurations for this scenario:

- [Example 8: Primary Unit Configuration, page B-21](#)
- [Example 8: Secondary Unit Configuration, page B-22](#)

## Example 8: Primary Unit Configuration

```

hostname pixfirewall
enable password myenablepassword

```

### Example 9: LAN-Based Active/Active Failover (Routed Mode)

```

password mypassword
interface Ethernet0
 nameif outside
 ip address 209.165.201.1 255.255.255.224 standby 209.165.201.2
 no shutdown
interface Ethernet1
 nameif inside
 ip address 192.168.2.1 255.255.255.0 standby 192.168.2.2
 no shutdown
interface Ethernet2
 description LAN Failover Interface
 no shutdown
interface ethernet3
 description STATE Failover Interface
telnet 192.168.2.45 255.255.255.255 inside
access-list acl_out permit tcp any host 209.165.201.5 eq 80
failover
failover lan unit primary
failover lan interface failover Ethernet2
failover lan enable
! The failover lan enable command is required on the PIX security appliance only.
failover polltime unit msec 200 holdtime msec 800
failover key key1
failover link state Ethernet3
failover interface ip failover 192.168.254.1 255.255.255.0 standby 192.168.254.2
failover interface ip state 192.168.253.1 255.255.255.0 standby 192.168.253.2
global (outside) 1 209.165.201.3 netmask 255.255.255.224
nat (inside) 1 0.0.0.0 0.0.0.0
static (inside,outside) 209.165.201.5 192.168.2.5 netmask 255.255.255.255 0 0
access-group acl_out in interface outside
route outside 0.0.0.0 0.0.0.0 209.165.201.4 1

```

## Example 8: Secondary Unit Configuration

```

failover
failover lan unit secondary
failover lan interface failover ethernet2
failover lan enable
failover key key1
failover interface ip failover 192.168.254.1 255.255.255.0 standby 192.168.254.2


```

## Example 9: LAN-Based Active/Active Failover (Routed Mode)

The following example shows how to configure Active/Active failover. In this example there are 2 user contexts, named admin and ctx1. [Figure B-8](#) shows the network diagram for the example.



Figure B-8 Active/Active Failover Configuration



See the following sections for the configurations for this scenario:

- [Example 9: Primary Unit Configuration](#)
- [Example 9: Secondary Unit Configuration](#)

## Example 9: Primary Unit Configuration

See the following sections for the primary unit configuration:

- [Example 9: Primary System Configuration, page B-23](#)
- [Example 9: Primary admin Context Configuration, page B-24](#)
- [Example 9: Primary ctx1 Context Configuration, page B-25](#)

## Example 9: Primary System Configuration

You must first enable multiple context mode using the **mode multiple** command. The mode is not stored in the configuration file, even though it endures reboots. Enter the **show mode** command to view the current mode.

```
hostname ciscopix
enable password farscape
password crichton
asdm image flash:/asdm.bin
```

### Example 9: LAN-Based Active/Active Failover (Routed Mode)

```

boot system flash:/cdisk.bin
mac-address auto
interface Ethernet0
 description LAN/STATE Failover Interface
interface Ethernet1
 no shutdown
interface Ethernet2
 no shutdown
interface Ethernet3
 no shutdown
interface Ethernet4
 no shutdown
interface Ethernet5
 no shutdown
interface Ethernet6
 no shutdown
interface Ethernet7
 no shutdown
interface Ethernet8
 no shutdown
interface Ethernet9
 no shutdown
failover
failover lan unit primary
failover lan interface folink Ethernet0
failover link folink Ethernet0
failover interface ip folink 10.0.4.1 255.255.255.0 standby 10.0.4.11
failover group 1
 primary
 preempt
failover group 2
 secondary
 preempt
admin-context admin
context admin
 description admin
 allocate-interface Ethernet1
 allocate-interface Ethernet2
 config-url flash:/admin.cfg
 join-failover-group 1
context ctx1
 description context 1
 allocate-interface Ethernet3
 allocate-interface Ethernet4
 config-url flash:/ctx1.cfg
 join-failover-group 2

```

### Example 9: Primary admin Context Configuration

```

enable password frek
password elixir
hostname admin
interface Ethernet1
 nameif outside
 security-level 0
 ip address 192.168.5.101 255.255.255.0 standby 192.168.5.111
interface Ethernet2
 nameif inside
 security-level 100
 ip address 192.168.0.1 255.255.255.0 standby 192.168.0.11
monitor-interface outside
monitor-interface inside

```

```
route outside 0.0.0.0 0.0.0.0 192.168.5.1 1
ssh 192.168.0.2 255.255.255.255 inside
```

## Example 9: Primary ctx1 Context Configuration

```
enable password quadrophenia
password tommy
hostname ctx1
interface Ethernet3
 nameif inside
 security-level 100
 ip address 192.168.20.1 255.255.255.0 standby 192.168.20.11
interface Ethernet4
 nameif outside
 security-level 0
 ip address 192.168.10.31 255.255.255.0 standby 192.168.10.41
 asr-group 1
access-list 201 extended permit ip any any
access-group 201 in interface outside
logging enable
logging console informational
monitor-interface inside
monitor-interface outside
route outside 0.0.0.0 0.0.0.0 192.168.10.71 1
```

## Example 9: Secondary Unit Configuration


You only need to configure the secondary security appliance to recognize the failover link. The secondary security appliance obtains the context configurations from the primary security appliance upon booting or when failover is first enabled. The **preempt** commands in the failover group configurations cause the failover groups to become active on their designated unit after the configurations have been synchronized and the preempt delay has passed.

```
failover
failover lan unit secondary
failover lan interface folink Ethernet0
failover interface ip folink 10.0.4.1 255.255.255.0 standby 10.0.4.11
```

## Example 10: Cable-Based Active/Standby Failover (Transparent Mode)

Figure B-6 shows the network diagram for a transparent mode failover configuration using a serial Failover cable. This configuration is only available on the PIX 500 series security appliance.

**Figure B-9** Transparent Mode Cable-Based Failover Configuration



The following are the typical commands in a cable-based, transparent firewall failover configuration.

```
enable password myenablepassword
passwd mypassword
hostname pixfirewall
asdm image flash:/asdm.bin
boot system flash:/image.bin
firewall transparent
interface Ethernet0
 speed 100
 duplex full
 nameif outside
 security-level 0
 no shutdown
interface Ethernet1
 speed 100
 duplex full
 nameif inside
 security-level 100
 no shutdown
interface Ethernet3
 description STATE Failover Interface
 telnet 192.168.2.45 255.255.255.255 mgmt
 access-list acl_in permit tcp any host 209.165.201.5 eq 80
 access-group acl_in in interface outside
```

```


ip address 209.165.201.1 255.255.255.0 standby 209.165.201.2
failover
failover link state Ethernet3
failover interface ip state 192.168.253.1 255.255.255.0 standby 192.168.253.2
route outside 0.0.0.0 0.0.0.0 209.165.201.4 1

```

## Example 11: LAN-Based Active/Standby Failover (Transparent Mode)

Figure B-7 shows the network diagram for a transparent mode failover configuration using an Ethernet failover link. The units are configured to detect unit failures and to fail over in under a second (see the **failover polltime unit** command in the primary unit configuration).

**Figure B-10** Transparent Mode LAN-Based Failover Configuration



See the following sections for the configurations for this scenario:

- [Example 8: Primary Unit Configuration, page B-21](#)
- [Example 8: Secondary Unit Configuration, page B-22](#)

## Example 11: Primary Unit Configuration

```

firewall transparent
hostname pixfirewall
enable password myenablepassword
password mypassword
interface Ethernet0
nameif outside

```

## Example 12: LAN-Based Active/Active Failover (Transparent Mode)

```

no shutdown
interface Ethernet1
 nameif inside
 no shutdown
interface Ethernet2
 description LAN Failover Interface
 no shutdown
interface ethernet3
 description STATE Failover Interface
telnet 192.168.2.45 255.255.255.255 inside
access-list acl_out permit tcp any host 209.165.201.5 eq 80
ip address 209.165.201.1 255.255.255.0 standby 209.165.201.2
failover
failover lan unit primary
failover lan interface failover Ethernet2
failover lan enable
! The failover lan enable command is required on the PIX security appliance only.
failover polltime unit msec 200 holdtime msec 800
failover key key1
failover link state Ethernet3
failover interface ip failover 192.168.254.1 255.255.255.0 standby 192.168.254.2
failover interface ip state 192.168.253.1 255.255.255.0 standby 192.168.253.2
access-group acl_out in interface outside
route outside 0.0.0.0 0.0.0.0 209.165.201.4 1

```

## Example 11: Secondary Unit Configuration

```


firewall transparent
failover
failover lan unit secondary
failover lan interface failover ethernet2
failover lan enable
failover key key1
failover interface ip failover 192.168.254.1 255.255.255.0 standby 192.168.254.2

```

## Example 12: LAN-Based Active/Active Failover (Transparent Mode)

The following example shows how to configure transparent mode Active/Active failover. In this example there are 2 user contexts, named admin and ctx1. [Figure B-8](#) shows the network diagram for the example.

**Figure B-11**     *Transparent Mode Active/Active Failover Configuration*



See the following sections for the configurations for this scenario:

- Example 9: Primary Unit Configuration
- Example 9: Secondary Unit Configuration

## Example 12: Primary Unit Configuration

See the following sections for the primary unit configuration:

- [Example 9: Primary System Configuration, page B-23](#)
- [Example 9: Primary admin Context Configuration, page B-24](#)
- [Example 9: Primary ctx1 Context Configuration, page B-25](#)

### Example 12: Primary System Configuration

You must first enable multiple context mode using the **mode multiple** command. The mode is not stored in the configuration file, even though it endures reboots. Enter the **show mode** command to view the current mode.

```
firewall transparent
```

## Example 12: LAN-Based Active/Active Failover (Transparent Mode)

```

hostname ciscopix
enable password farscape
password crichton
asdm image flash:/asdm.bin
boot system flash:/cdisk.bin
mac-address auto
interface Ethernet0
 description LAN/STATE Failover Interface
interface Ethernet1
 no shutdown
interface Ethernet2
 no shutdown
interface Ethernet3
 no shutdown
interface Ethernet4
 no shutdown
interface Ethernet5
 no shutdown
interface Ethernet6
 no shutdown
interface Ethernet7
 no shutdown
interface Ethernet8
 no shutdown
interface Ethernet9
 no shutdown
failover
failover lan unit primary
failover lan interface folink Ethernet0
failover link folink Ethernet0
failover interface ip folink 10.0.4.1 255.255.255.0 standby 10.0.4.11
failover group 1
 primary
 preempt
failover group 2
 secondary
 preempt
admin-context admin
context admin
 description admin
 allocate-interface Ethernet1
 allocate-interface Ethernet2
 config-url flash:/admin.cfg
 join-failover-group 1
context ctx1
 description context 1
 allocate-interface Ethernet3
 allocate-interface Ethernet4
 config-url flash:/ctx1.cfg
 join-failover-group 2

```

## Example 12: Primary admin Context Configuration

```

enable password frek
password elixir
hostname admin
interface Ethernet1
 nameif outside
 security-level 0
interface Ethernet2
 nameif inside
 security-level 100

```



```
ip address 192.168.5.31 255.255.255.0 standby 192.168.5.32
monitor-interface outside
monitor-interface inside
route outside 0.0.0.0 0.0.0.0 192.168.5.1 1
ssh 192.168.5.72 255.255.255.255 inside
```

## Example 12: Primary ctx1 Context Configuration

```
enable password quadrophenia
password tommy
hostname ctx1
interface Ethernet3
 nameif inside
 security-level 100
interface Ethernet4
 nameif outside
 security-level 0
access-list 201 extended permit ip any any
access-group 201 in interface outside
logging enable
logging console informational
ip address 192.168.10.31 255.255.255.0 standby 192.168.10.32
monitor-interface inside
monitor-interface outside
route outside 0.0.0.0 0.0.0.0 192.168.10.1 1
```

## Example 12: Secondary Unit Configuration

You only need to configure the secondary security appliance to recognize the failover link. The secondary security appliance obtains the context configurations from the primary security appliance upon booting or when failover is first enabled. The **preempt** commands in the failover group configurations cause the failover groups to become active on their designated unit after the configurations have been synchronized and the preempt delay has passed.


```
firewall transparent
failover
failover lan unit secondary
failover lan interface folink Ethernet0
failover interface ip folink 10.0.4.1 255.255.255.0 standby 10.0.4.11
```

## Example 13: Dual ISP Support Using Static Route Tracking

This configuration shows a remote office using static route tracking to use a backup ISP route if the primary ISP route fails. The security appliance in the remote office uses ICMP echo requests to monitor the availability of the main office gateway. If that gateway becomes unavailable through the default route, the default route is removed from the routing table and the floating route to the backup ISP is used in its place.

## Example 13: Dual ISP Support Using Static Route Tracking

Figure B-12 Dual ISP Support



```

passwd password1
enable password password2
hostname myfirewall
asdm image disk0:/asdm.bin
boot system disk0:/image.bin
!
interface gigabitethernet 0/0
 nameif outside
 security-level 0
 ip address 10.1.1.2 255.255.255.0
 no shutdown
!
interface gigabitethernet 0/1
 description backup isp link
 nameif backupisp
 security-level 100
 ip address 172.16.2.2 255.255.255.0
 no shutdown
!
sla monitor 123
 type echo protocol ipIcmpEcho 10.2.1.2 interface outside
 num-packets 3
 timeout 1000
 frequency 3
sla monitor schedule 123 life forever start-time now
!
track 1 rtr 123 reachability
!
route outside 0.0.0.0 0.0.0.0 10.1.1.1 track 1
! The above route is used while the tracked object, router 10.2.1.2
! is available. It is removed when the router becomes unavailable.
!
route backupisp 0.0.0.0 0.0.0.0 172.16.2.1 254
! The above route is a floating static route that is added to the
! routing table when the tracked route is removed.


```

153924

## Example 14: ASA 5505 Base License

This configuration creates three VLANs: inside (business), outside (Internet), and home (see [Figure B-13](#)). Both the home and inside VLANs can access the outside, but the home VLAN cannot access the inside VLAN. The inside VLAN can access the home VLAN so both VLANs can share a printer. Because the outside IP address is set using DHCP, the inside and home VLANs use interface PAT when accessing the Internet.

**Figure B-13** Example 13



```
passwd g00fba11
enable password genlu$
hostname Buster
asdm image disk0:/asdm.bin
boot system disk0:/image.bin
interface vlan 2
 nameif outside
 security-level 0
 ip address dhcp setroute
 no shutdown
interface vlan 1
 nameif inside
 security-level 100
 ip address 192.168.1.1 255.255.255.0
 no shutdown
interface vlan 3
 ! This interface cannot communicate with the inside interface. This is required using
 ! the Base license
 no forward interface vlan 1
 nameif home
 security-level 50
 ip address 192.168.2.1 255.255.255.0
 no shutdown
interface ethernet 0/0
 switchport access vlan 2
 no shutdown
interface ethernet 0/1
```

## Example 14: ASA 5505 Base License

```


 switchport access vlan 1
 no shutdown
interface ethernet 0/2
 switchport access vlan 1
 no shutdown
interface ethernet 0/3
 switchport access vlan 3
 no shutdown
interface ethernet 0/4
 switchport access vlan 3
 no shutdown
interface ethernet 0/5
 switchport access vlan 3
 no shutdown
interface ethernet 0/6
 description PoE for IP phone1
 switchport access vlan 1
 no shutdown
interface ethernet 0/7
 description PoE for IP phone2
 switchport access vlan 1
 no shutdown
nat (inside) 1 0 0
nat (home) 1 0 0
global (outside) 1 interface
! The previous NAT statements match all addresses on inside and home, so you need to
! also perform NAT when hosts access the inside or home networks (as well as the outside).
! Or you can exempt hosts from NAT for inside <--> home traffic, as effected by the
! following:
access-list natexempt-inside extended permit ip any 192.168.2.0 255.255.255.0
access-list natexempt-home extended permit ip any 192.168.1.0 255.255.255.0
nat (inside) 0 access-list natexempt-inside
nat (home) 0 access-list natexempt-home
http server enable
http 192.168.1.0 255.255.255.0 inside
dhcpd address 192.168.1.2-192.168.1.254 inside
dhcpd auto_config outside
dhcpd enable inside
logging asdm informational
ssh 192.168.1.0 255.255.255.0 inside

```

## Example 15: ASA 5505 Security Plus License with Failover and Dual-ISP Backup

This configuration creates five VLANs: inside, outside, dmz, backup-isp and faillink (see [Figure B-13](#)).

**Figure B-14** Example 15



See the following sections for the configurations for this scenario:

- [Example 15: Primary Unit Configuration](#)
- [Example 15: Secondary Unit Configuration](#)

### Example 15: Primary Unit Configuration

```
passwd g00fba11
enable password gen1u$
```

## Example 15: ASA 5505 Security Plus License with Failover and Dual-ISP Backup

```

hostname Buster
asdm image disk0:/asdm.bin
boot system disk0:/image.bin
interface vlan 2
 description Primary ISP interface
 nameif outside
 security-level 0
 ip address 209.165.200.224 standby 209.165.200.225
 backup interface vlan 4
 no shutdown
interface vlan 1
 nameif inside
 security-level 100
 ip address 192.168.1.1 255.255.255.0
 no shutdown
interface vlan 3
 nameif dmz
 security-level 50
 ip address 192.168.2.1 255.255.255.0
 no shutdown
interface vlan 4
 description Backup ISP interface
 nameif backup-isp
 security-level 0
 ip address 209.168.202.128 standby 209.168.202.129
 no shutdown
interface vlan 5
 description LAN Failover Interface
interface ethernet 0/0
 switchport access vlan 2
 no shutdown
interface ethernet 0/1
 switchport access vlan 4
 no shutdown
interface ethernet 0/2
 switchport access vlan 1
 no shutdown
interface ethernet 0/3
 switchport access vlan 3
 no shutdown
interface ethernet 0/4
 switchport access vlan 5
 no shutdown
failover
failover lan unit primary
failover lan interface faillink vlan5
failover lan faillink vlan5
failover polltime unit 3 holdtime 10
failover key key1
failover interface ip faillink 10.1.1.1 255.255.255.0 standby 10.1.1.2
nat (inside) 1 0 0
nat (home) 1 0 0
global (outside) 1 interface
! The previous NAT statements match all addresses on inside and home, so you need to
! also perform NAT when hosts access the inside or home networks (as well as the outside).
! Or you can exempt hosts from NAT for inside <--> home traffic, as effected by the
! following:
access-list natexempt-inside extended permit ip any 192.168.2.0 255.255.255.0
access-list natexempt-home extended permit ip any 192.168.1.0 255.255.255.0
nat (inside) 0 access-list natexempt-inside
nat (home) 0 access-list natexempt-home
sla monitor 123
 type echo protocol ipIcmpEcho 209.165.200.234 interface outside
 num-packets 2

```

```

frequency 5
sla monitor schedule 123 life forever start-time now
track 1 rtr 123 reachability
route outside 0 0 209.165.200.234 1 track 1
! This route is for the primary ISP.
route backup-isp 0 0 209.165.202.154 2
! If the link goes down for the primary ISP, either due to a hardware failure
! or unplugged cable, then this route will be used.
http server enable
http 192.168.1.0 255.255.255.0 inside
dhcpd address 192.168.1.2-192.168.1.254 inside
dhcpd auto_config outside
dhcpd enable inside
logging asdm informational
ssh 192.168.1.0 255.255.255.0 inside

```

## Example 15: Secondary Unit Configuration

You only need to configure the secondary security appliance to recognize the failover link. The secondary security appliance obtains the context configurations from the primary security appliance upon booting or when failover is first enabled.

```

interface ethernet 0/4
 switchport access vlan 5
 no shutdown
failover
failover lan unit secondary
failover lan interface faillink vlan5
failover polltime unit 3 holdtime 10
failover key key1
failover interface ip faillink 10.1.1.1 255.255.255.0 standby 10.1.1.2

```

## Example 16: Network Traffic Diversion

The following configuration example shows the ASA 5500 series adaptive security appliance with Version 7.2.1 software and the AIP SSM module with IPS software 5.1.1.

Network traffic that traverses the adaptive security appliance includes internal users who access the Internet, Internet users who access resources protected by an adaptive security appliance in a demilitarized zone (DMZ), or in an inside network. Network traffic sent to and from the adaptive security appliance is not sent to the IPS module for inspection. Examples of traffic not sent to the IPS module include pinging (through ICMP) of the adaptive security appliance interfaces or Telnetting to the adaptive security appliance.

The required configuration components for the ASA 5510 adaptive security appliance include interfaces, access lists, network address translation (NAT), and routing. The required configuration components for the AIP SSM include the network setup, allowed hosts, interface configuration, signature definitions, and event action rules.

To obtain more information about the commands used in this section, use the Command Lookup Tool (for registered customers only).

**Note**

The IP addressing schemes used in this configuration are not legally routable on the Internet. These schemes are RFC 1918 addresses that have been used in a test environment.

Figure B-15 shows the network diagram for this configuration example.

**Figure B-15 Network Diagram**




Figure B-16 on page B-39 and Figure B-17 on page B-41 show the initial configurations for the ASA 5510 adaptive security appliance and AIP SSM.



**Figure B-16 Configuration for the ASA 5510 Adaptive Security Appliance**

```

onionlabaip#show configuration
! -----
! Version 5.1(1)
! Current configuration last modified Wed Aug 23 09:26:03 2006
! -----
service interface
exit
! -----
service analysis-engine
virtual-sensor vs0
physical-interface GigabitEthernet0/1
exit
exit
! -----
service authentication
exit
! -----
service event-action-rules rules0

!--- The variables are defined.

variables DMZ address 192.168.1.0-192.168.1.255
variables IN address 10.2.2.0-10.2.2.255
exit
! -----
service host
network-settings

!--- The management IP address is set.

host-ip 172.22.1.169/24,172.22.1.1
host-name onionlabaip
telnet-option disabled
access-list x.x.0.0/16

!--- The access list IP address is removed from the configuration
!--- because the specific IP address is not relevant to this document.

exit
time-zone-settings
offset -360
standard-time-zone-name GMT-06:00
exit
summertime-option recurring
offset 60
summertime-zone-name UTC
start-summertime
month april
week-of-month first
day-of-week sunday
time-of-day 02:00:00
exit
end-summertime
month october
week-of-month last
day-of-week sunday
time-of-day 02:00:00
exit
exit

asdm image disk0:/asdm521.bin
no asdm history enable
arp timeout 14400

```

191240

## Example 16: Network Traffic Diversion

```

!--- Translation rules are added.

global (outside) 1 172.16.1.100
global (dmz) 1 192.168.1.100
nat (inside) 1 10.2.2.0 255.255.255.0
static (dmz,outside) 172.16.1.50 192.168.1.50 netmask 255.255.255.255
static (inside,dmz) 10.2.2.200 10.2.2.200 netmask 255.255.255.255

!--- Access lists are applied to the interfaces.

access-group acl_outside_in in interface outside
access-group acl_inside_in in interface inside
access-group acl_dmz_in in interface dmz
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat 0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect 0:02:00
timeout uauth 0:05:00 absolute
no snmp-server location
no snmp-server contact
snmp-server enable traps snmp authentication linkup linkdown coldstart
telnet timeout 5
ssh timeout 5
console timeout 0
!
class-map inspection_default

 match default-inspection-traffic
!
!
policy-map type inspect dns preset_dns_map
parameters
 message-length maximum 512
policy-map global_policy

!--- Out-of-the-box default configuration includes
!--- policy-map global_policy.

class inspection_default
 inspect dns preset_dns_map
 inspect ftp
 inspect h323 h225
 inspect h323 ras
 inspect netbios
 inspect rsh
 inspect rtsp
 inspect skinny
 inspect esmtp
 inspect sqlnet
 inspect sunrpc
 inspect tftp
 inspect sip
 inspect xdmcp
!
service-policy global_policy global

!--- Out-of-the-box default configuration includes
!--- the service-policy global_policy applied globally.

prompt hostname context
.
: end

```

**Figure B-17 Configuration for the AIP SSM**

```

ciscoasa#show running-config
: Saved
:
ASA Version 7.2(1)
!
hostname ciscoasa
enable password WwXYvtKrnjXqGbu1 encrypted
names
!

!--- IP addressing is added to the default configuration.

interface Ethernet0/0
 nameif outside
 security-level 0
 ip address 172.16.1.254 255.255.255.0
!
interface Ethernet0/1
 nameif inside
 security-level 100
 ip address 10.2.2.254 255.255.255.0
!
interface Ethernet0/2
 nameif dmz
 security-level 50
 ip address 192.168.1.254 255.255.255.0
!
interface Management0/0
 nameif management
 security-level 0
 ip address 172.22.1.160 255.255.255.0
 management-only
!
passwd 2KFQnbNIdI.2KYOU encrypted
ftp mode passive

!--- Access lists are added in order to allow test
!--- traffic (ICMP and Telnet).

access-list acl_outside_in extended permit icmp any host 172.16.1.50
access-list acl_inside_in extended permit ip 10.2.2.0 255.255.255.0 any
access-list acl_dmz_in extended permit icmp 192.168.1.0 255.255.255.0 any
pager lines 24

!--- Logging is enabled.

logging enable
logging buffered debugging
mtu outside 1500
mtu inside 1500
mtu dmz 1500
mtu management 1500

exit
! -----
service logger
exit
! -----

```

19/241

## Example 16: Network Traffic Diversion

```

service network-access
exit
! -----
service notification
exit
! -----
service signature-definition sig0

!--- The signature is modified from the default setting for testing purposes.

signatures 2000 0
alert-severity high
engine atomic-ip
event-action produce-alert|produce-verbose-alert
exit
alert-frequency
summary-mode fire-all
summary-key AxBx
exit
exit
status
enabled true
exit
exit

!--- The signature is modified from the default setting for testing purposes.

signatures 2004 0
alert-severity high
engine atomic-ip
event-action produce-alert|produce-verbose-alert
exit
alert-frequency
summary-mode fire-all
summary-key AxBx
exit
exit
status
enabled true
exit
exit

!--- The custom signature is added for testing purposes.

signatures 60000 0
alert-severity high
sig-fidelity-rating 75
sig-description
sig-name Telnet Command Authorization Failure
sig-string-info Command authorization failed
sig-comment signature triggers string command authorization failed
exit
engine atomic-ip
specify-l4-protocol yes
l4-protocol tcp
no tcp-flags
no tcp-mask
exit
specify-payload-inspection yes

regex-string Command authorization failed
exit
exit
exit

```

```

exit
exit
! -----
service ssh-known-hosts
exit
! -----
service trusted-certificates
exit
! -----
service web-server
enable-tls true
exit
onionlabaip#

```

## Inspecting All Traffic with the AIP SSM

This configuration meets the requirement to monitor all traffic. In addition, you must make two decisions about how the ASA 5510 and AIP SSM interact.

- Is the AIP SSM module to be deployed in promiscuous or inline mode?

Promiscuous mode means that a copy of the data is sent to the AIP SSM while the ASA 5510 forwards the original data to the destination. The AIP SSM in promiscuous mode can be considered as an intrusion detection system (IDS). In this mode, the trigger packet that causes the alarm can still reach the destination. Shunning can occur and stop additional packets from reaching the destination; however, the trigger packet is not stopped.

Inline mode means that the ASA 5510 forwards the data to the AIP SSM for inspection. If the data passes AIP SSM inspection, the data returns to the ASA 5510 in order to continue being processed and sent to the destination. The AIP SSM in inline mode can be considered to be an intrusion prevention system (IPS). Unlike promiscuous mode, an inline mode IPS can actually stop the trigger packet from reaching the destination.

- If the ASA 5510 cannot communicate with the AIP SSM, how should the adaptive security appliance handle traffic for inspection?

Examples of instances when the ASA 5510 cannot communicate with the AIP SSM include AIP SSM reloads or whether the module fails and needs replacement. In this case, the adaptive security appliance can fail-open or fail-closed.

Fail-open allows the adaptive security appliance to continue to pass traffic for inspection to the final destination if the AIP SSM cannot be reached. Fail-closed blocks traffic for inspection when the adaptive security appliance cannot communicate with the AIP SSM.



### Note

Define the traffic for inspection with an access list. In the following example, the access list permits all IP traffic from any source to any destination. Therefore, traffic for inspection can be anything that passes through the adaptive security appliance.

```

ciscoasa(config)#access-list traffic_for_ips permit ip any any
ciscoasa(config)#class-map ips_class_map
ciscoasa(config-cmap)#match access-list traffic_for_ips

!--- The match any
!--- command can be used in place of the match access-list [access-list name]
!--- command. In this example, access-list traffic_for_ips permits
!--- all traffic. The match any command also

```

```

!--- permits all traffic. You can use either configuration.
!--- When you define an access-list, it can ease troubleshooting.

ciscoasa(config)#policy-map global_policy

!--- Note that policy-map global_policy is a part of the
!--- default configuration. In addition, policy-map global_policy is applied
!--- globally using the service-policy command.

ciscoasa(config-pmap)#class ips_class_map
ciscoasa(config-pmap-c)#ips inline fail-open

!--- Two decisions need to be made.
!--- First, does the AIP-SSM function
!--- in inline or promiscuous mode?
!--- Second, does the ASA fail-open or fail-closed?

```

## Inspecting Specific Traffic with the AIP SSM

If you want the AIP SSM to monitor a subset of all traffic, you can modify two independent variables on the adaptive security appliance:

- You can write the access list to include or exclude the necessary traffic.
- You can apply a service policy to an interface or globally.

The network diagram in [Figure B-15](#) shows the AIP SSM inspecting all traffic between the outside network and the DMZ network, as shown in the following example:

```

ciscoasa#configure terminal
ciscoasa(config)#access-list traffic_for_ips deny ip 10.2.2.0 255.255.255.0 192.168.1.0
255.255.255.0
ciscoasa(config)#access-list traffic_for_ips permit ip any 192.168.1.0 255.255.255.0
ciscoasa(config)#access-list traffic_for_ips deny ip 192.168.1.0 255.255.255.0 10.2.2.0
255.255.255.0
ciscoasa(config)#access-list traffic_for_ips permit ip 192.168.1.0 255.255.255.0 any
ciscoasa(config)#class-map ips_class_map
ciscoasa(config-cmap)#match access-list traffic_for_ips
ciscoasa(config)#policy-map interface_policy
ciscoasa(config-pmap)#class ips_class_map
ciscoasa(config-pmap-c)#ips inline fail-open
ciscoasa(config)#service-policy interface_policy interface dmz

!--- The access-list denies traffic from the inside network to the DMZ network
!--- and traffic to the inside network from the DMZ network.
!--- In addition, the service-policy command is applied to the DMZ interface.

```

The following example shows how to configure the AIP SSM to monitor traffic from the inside network to the outside network, but exclude the inside network to the DMZ network.



### Note

You must have an intermediate understanding of statefulness, TCP, UDP, ICMP, connection, and connectionless communications to understand the following example.

```

ciscoasa#configure terminal
ciscoasa(config)#access-list traffic_for_ips deny ip 10.2.2.0 255.255.255.0
192.168.1.0 255.255.255.0
ciscoasa(config)#access-list traffic_for_ips permit ip 10.2.2.0
255.255.255.0 any
ciscoasa(config)#class-map ips_class_map

```

```

ciscoasa(config-cmap)#match access-list traffic_for_ips
ciscoasa(config)#policy-map interface_policy
ciscoasa(config-pmap)#class ips_class_map
ciscoasa(config-pmap-c)#ips inline fail-open
ciscoasa(config)#service-policy interface_policy interface inside

```

The access list denies traffic initiated on the inside network destined for the DMZ network. The second access list line permits or sends traffic initiated on the inside network destined for the outside network to the AIP SSM. At this point the statefulness of the adaptive security appliance comes into play.

For example, an internal user initiates a TCP connection (Telnet) to a device on the outside network (router). The user successfully connects to the router and logs in, then issues a router command that is not authorized. The router responds with the message, “Command authorization failed.” The data packet that contains the message, “Command authorization failed” has the outside router as the source and the inside user as the destination. The source (outside) and destination (inside) do not match the access lists previously defined. The adaptive security appliance keeps track of stateful connections. As a result, the returning data packet (outside to inside) is sent to the AIP SSM for inspection. Custom signature 60000 0 (configured on the AIP SSM) alarms.

**Note**

By default, the adaptive security appliance does not maintain state for the ICMP traffic. In the previous example, the internal user pings (ICMP echo request) the outside router. The router responds with an ICMP echo-reply. The AIP SSM inspects the echo request packet, but not the echo-reply packet. If ICMP inspection is enabled on the adaptive security appliance, both the echo request and echo-reply packets are inspected by the AIP SSM.

## Verifying the Recording of Alert Events

To verify that alert events are recorded in the AIP SSM, perform the following steps:

**Step 1** Log into the AIP SSM with the administrator user account.

**Note**

The output varies according to signature settings, the type of traffic sent to the AIP SSM, and network load.

The Output Interpreter Tool (OIT), for registered customers only, supports certain **show** commands. Use the OIT to view an analysis of **show** command output. This tool is one of a set of support tools, available at <http://www.cisco.com/public/support/tac/tools.shtml>.

**Step 2** Enter the **show events alert** command.

The following output appears.

```

evIdsAlert: eventId=1156198930427770356 severity=high vendor=Cisco
 originator:
 hostId: onionlabaip
 appName: sensorApp
 appInstanceId: 345
time: 2006/08/24 18:52:57 2006/08/24 13:52:57 UTC
signature: description=Telnet Command Authorization Failure id=60000 version=custom
 subsigId: 0
 sigDetails: Command authorization failed
interfaceGroup:
vlan: 0
participants:

```

## Example 16: Network Traffic Diversion

```

attacker:
 addr: locality=OUT 172.16.1.200
 port: 23
target:
 addr: locality=IN 10.2.2.200
 port: 33189
riskRatingValue: 75
interface: ge0_1
protocol: tcp

evIdsAlert: eventId=1156205750427770078 severity=high vendor=Cisco
 originator:
 hostId: onionlabaip
 appName: sensorApp
 appInstanceId: 345
time: 2006/08/24 19:46:08 2006/08/24 14:46:08 UTC
signature: description=ICMP Echo Request id=2004 version=S1
 subsigId: 0
interfaceGroup:
vlan: 0
participants:
 attacker:
 addr: locality=OUT 172.16.1.200
 target:
 addr: locality=DMZ 192.168.1.50
triggerPacket:
000000 00 16 C7 9F 74 8C 00 15 2B 95 F9 5E 08 00 45 00 t...+...^..E.
000010 00 3C 2A 57 00 00 FF 01 21 B7 AC 10 01 C8 C0 A8 .<*W....!.....
000020 01 32 08 00 F5 DA 11 24 00 00 00 01 02 03 04 05 .2.....$.
000030 06 07 08 09 0A 0B 0C 0D 0E 0F 10 11 12 13 14 15
000040 16 17 18 19 1A 1B 1C 1D 1E 1F
riskRatingValue: 100
interface: ge0_1
protocol: icmp

evIdsAlert: eventId=1156205750427770079 severity=high vendor=Cisco originator:
hostId: onionlabaip
appName: sensorApp
appInstanceId: 345
time: 2006/08/24 19:46:08 2006/08/24 14:46:08 UTC
signature: description=ICMP Echo Reply id=2000 version=S1
 subsigId: 0
interfaceGroup:
vlan: 0
participants:
 attacker:
 addr: locality=DMZ 192.168.1.50
 target:
 addr: locality=OUT 172.16.1.200
triggerPacket:
000000 00 16 C7 9F 74 8E 00 03 E3 02 6A 21 08 00 45 00 t.....j!..E.
000010 00 3C 2A 57 00 00 FF 01 36 4F AC 10 01 32 AC 10 .<*W....6O...2..
000020 01 C8 00 00 FD DA 11 24 00 00 00 01 02 03 04 05$.
000030 06 07 08 09 0A 0B 0C 0D 0E 0F 10 11 12 13 14 15
000040 16 17 18 19 1A 1B 1C 1D 1E 1F
riskRatingValue: 100
interface: ge0_1
protocol: icmp

```



In these configurations, several IPS signatures are tuned to alarm on test traffic. Signatures 2000 and 2004 are modified. Custom signature 60000 is added. In a network where little data passes through the adaptive security appliance, you may need to modify signatures in order to trigger events. If the adaptive security appliance and AIP SSM are deployed in an environment that passes a large amount of traffic, the default signature settings will probably generate an event.

## Troubleshooting the Configuration

To troubleshoot your configuration, perform the following steps:

The OIT (for registered customers only) supports certain **show** commands. Use the OIT to view an analysis of **show** command output.

**Step 1** From the ASA 5510, enter these **show** commands:

- a. **show module**—Shows information about the SSM on the adaptive security appliance as well as system information.

```
ciscoasa#show module
Mod Card Type Model Serial No.

0 ASA 5510 Adaptive Security Appliance ASA5510 JMX1016K0RN
1 ASA 5500 Series Security Services Module-10 ASA-SSM-10 JAB101502A6

Mod MAC Address Range Hw Version Fw Version Sw Version

0 0016.c79f.748c to 0016.c79f.7490 1.1 1.0(10)0 7.2(1)
1 0016.c79f.7567 to 0016.c79f.7567 1.0 1.0(10)0 5.1(1)S205.0

Mod SSM Application Name Status SSM Application Version

1 IPS Up 5.1(1)S205.0

Mod Status Data Plane Status Compatibility

0 Up Sys Not Applicable
1 Up Up

!--- Each of the areas highlighted indicate that
!--- the ASA recognizes the AIP-SSM and the AIP-SSM status is up.
```

- b. **show run**—Shows the current running configuration on the adaptive security appliance.

```
ciscoasa#show run

!--- Output is suppressed.

access-list traffic_for_ips extended permit ip any any
...
class-map ips_class_map
match access-list traffic_for_ips
...
policy-map global_policy
...
class ips_class_map
ips inline fail-open
...
service-policy global_policy global
```

## ■ Example 16: Network Traffic Diversion

```
!--- Each of these lines are needed
!--- in order to send data to the AIP-SSM.
```

c. □ **show access-list** —Shows the counters for an access list.

```
ciscoasa#show access-list traffic_for_ips
access-list traffic_for_ips; 1 elements
access-list traffic_for_ips line 1 extended permit ip any any (hitcnt=2) 0x9bea7286

!--- Confirms the access-list displays a hit count greater than zero.
```

**Step 2** Before you install and use the AIP SSM, if network traffic does not pass through the adaptive security appliance as expected, troubleshoot the network and ASA 5510 access policy rules.

---



# APPENDIX C

## Using the Command-Line Interface

---

This appendix describes how to use the CLI on the security appliance, and includes the following sections:

- [Firewall Mode and Security Context Mode, page C-1](#)
- [Command Modes and Prompts, page C-2](#)
- [Syntax Formatting, page C-3](#)
- [Abbreviating Commands, page C-3](#)
- [Command-Line Editing, page C-3](#)
- [Command Completion, page C-4](#)
- [Command Help, page C-4](#)
- [Filtering show Command Output, page C-4](#)
- [Command Output Paging, page C-5](#)
- [Adding Comments, page C-6](#)
- [Text Configuration Files, page C-6](#)



### Note

The CLI uses similar syntax and other conventions to the Cisco IOS CLI, but the security appliance operating system is not a version of Cisco IOS software. Do not assume that a Cisco IOS CLI command works with or has the same function on the security appliance.

---

## Firewall Mode and Security Context Mode

The security appliance runs in a combination of the following modes:

- Transparent firewall or routed firewall mode

The firewall mode determines if the security appliance runs as a Layer 2 or Layer 3 firewall.

- Multiple context or single context mode

The security context mode determines if the security appliance runs as a single device or as multiple security contexts, which act like virtual devices.

Some commands are only available in certain modes.

# Command Modes and Prompts

The security appliance CLI includes command modes. Some commands can only be entered in certain modes. For example, to enter commands that show sensitive information, you need to enter a password and enter a more privileged mode. Then, to ensure that configuration changes are not entered accidentally, you have to enter a configuration mode. All lower commands can be entered in higher modes, for example, you can enter a privileged EXEC command in global configuration mode.

**Note**

The various types of prompts are all default prompts and when configured, they can be different.

- When you are in the system configuration or in single context mode, the prompt begins with the hostname:  
`hostname`
- When printing the prompt string, the prompt configuration is parsed and the configured keyword values are printed in the order in which you have set the **prompt** command. The keyword arguments can be any of the following and in any order: hostname, domain, context, priority, state.  
`asa(config)# prompt hostname context priority state`
- When you are within a context, the prompt begins with the hostname followed by the context name:  
`hostname/context`

The prompt changes depending on the access mode:

- User EXEC mode  
User EXEC mode lets you see minimum security appliance settings. The user EXEC mode prompt appears as follows when you first access the security appliance:  
`hostname>`  
`hostname/context>`
- Privileged EXEC mode  
Privileged EXEC mode lets you see all current settings up to your privilege level. Any user EXEC mode command will work in privileged EXEC mode. Enter the **enable** command in user EXEC mode, which requires a password, to start privileged EXEC mode. The prompt includes the number sign (#):  
`hostname#`  
`hostname/context#`
- Global configuration mode  
Global configuration mode lets you change the security appliance configuration. All user EXEC, privileged EXEC, and global configuration commands are available in this mode. Enter the **configure terminal** command in privileged EXEC mode to start global configuration mode. The prompt changes to the following:  
`hostname(config)#`  
`hostname/context(config)#`
- Command-specific configuration modes

From global configuration mode, some commands enter a command-specific configuration mode. All user EXEC, privileged EXEC, global configuration, and command-specific configuration commands are available in this mode. For example, the **interface** command enters interface configuration mode. The prompt changes to the following:

```
hostname(config-if)#

hostname/context(config-if)#
```

## Syntax Formatting

Command syntax descriptions use the following conventions:

**Table C-1**      **Syntax Conventions**

| Convention     | Description                                                                                                                                                                                                                  |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>bold</b>    | Bold text indicates commands and keywords that you enter literally as shown.                                                                                                                                                 |
| <i>italics</i> | Italic text indicates arguments for which you supply values.                                                                                                                                                                 |
| [x]            | Square brackets enclose an optional element (keyword or argument).                                                                                                                                                           |
|                | A vertical bar indicates a choice within an optional or required set of keywords or arguments.                                                                                                                               |
| [x   y]        | Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.                                                                                                                     |
| {x   y}        | Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.                                                                                                                               |
| [x {y   z}]    | Nested sets of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element. |

## Abbreviating Commands

You can abbreviate most commands down to the fewest unique characters for a command; for example, you can enter **wr t** to view the configuration instead of entering the full command **write terminal**, or you can enter **en** to start privileged mode and **conf t** to start configuration mode. In addition, you can enter **o** to represent **o.o.o.o.**

## Command-Line Editing

The security appliance uses the same command-line editing conventions as Cisco IOS software. You can view all previously entered commands with the **show history** command or individually with the up arrow or **^p** command. Once you have examined a previously entered command, you can move forward in the list with the down arrow or **^n** command. When you reach a command you wish to reuse, you can edit it or press the **Enter** key to start it. You can also delete the word to the left of the cursor with **^w**, or erase the line with **^u**.

The security appliance permits up to 512 characters in a command; additional characters are ignored.

# Command Completion

To complete a command or keyword after entering a partial string, press the **Tab** key. The security appliance only completes the command or keyword if the partial string matches only one command or keyword. For example, if you enter **s** and press the **Tab** key, the security appliance does not complete the command because it matches more than one command. However, if you enter **dis**, the **Tab** key completes the command **disable**.

## Command Help

Help information is available from the command line by entering the following commands:

- **help** *command\_name*  
Shows help for the specific command.
- *command\_name* **?**  
Shows a list of arguments available.
- *string?* (no space)  
Lists the possible commands that start with the string.
- **?** and **+?**  
Lists all commands available. If you enter **?**, the security appliance shows only commands available for the current mode. To show all commands available, including those for lower modes, enter **+?**.

**Note**

If you want to include a question mark (?) in a command string, you must press **Ctrl-V** before typing the question mark so you do not inadvertently invoke CLI help.

## Filtering show Command Output

You can use the vertical bar (|) with any **show** command and include a filter option and filtering expression. The filtering is performed by matching each output line with a regular expression, similar to Cisco IOS software. By selecting different filter options you can include or exclude all output that matches the expression. You can also display all output beginning with the line that matches the expression.

The syntax for using filtering options with the **show** command is as follows:

```
hostname# show command | {include | exclude | begin | grep [-v]} regex
```

In this command string, the first vertical bar (|) is the operator and must be included in the command. This operator directs the output of the **show** command to the filter. In the syntax diagram, the other vertical bars (|) indicate alternative options and are not part of the command.

The **include** option includes all output lines that match the regular expression. The **grep** option without **-v** has the same effect. The **exclude** option excludes all output lines that match the regular expression. The **grep** option with **-v** has the same effect. The **begin** option shows all the output lines starting with the line that matches the regular expression.

Replace *regexp* with any Cisco IOS regular expression. See The regular expression is not enclosed in quotes or double-quotes, so be careful with trailing white spaces, which will be taken as part of the regular expression.

When creating regular expressions, you can use any letter or number that you want to match. In addition, certain keyboard characters have special meaning when used in regular expressions. Table C-2 lists the keyboard characters that have special meaning.

**Table C-2** Using Special Characters in Regular Expressions

| Character Type | Character      | Special Meaning                                                                                                                                                        |
|----------------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| period         | .              | Matches any single character, including white space.                                                                                                                   |
| asterisk       | *              | Matches 0 or more sequences of the pattern.                                                                                                                            |
| plus sign      | +              | Matches 1 or more sequences of the pattern.                                                                                                                            |
| question mark  | ? <sup>1</sup> | Matches 0 or 1 occurrences of the pattern.                                                                                                                             |
| caret          | ^              | Matches the beginning of the input string.                                                                                                                             |
| dollar sign    | \$             | Matches the end of the input string.                                                                                                                                   |
| underscore     | _              | Matches a comma (,), left brace ({), right brace (}), left parenthesis, right parenthesis, the beginning of the input string, the end of the input string, or a space. |
| brackets       | []             | Designates a range of single-character patterns.                                                                                                                       |
| hyphen         | -              | Separates the end points of a range.                                                                                                                                   |

1. Precede the question mark with **Ctrl-V** to prevent the question mark from being interpreted as a help command.

To use these special characters as single-character patterns, remove the special meaning by preceding each character with a backslash (\).

## Command Output Paging

On commands such as **help** or **?**, **show**, **show xlate**, or other commands that provide long listings, you can determine if the information displays a screen and pauses, or lets the command run to completion. The **pager** command lets you choose the number of lines to display before the More prompt appears.

When paging is enabled, the following prompt appears:

```
<--- More --->
```

The More prompt uses syntax similar to the UNIX **more** command:

- To view another screen, press the Space bar.
- To view the next line, press the **Enter** key.
- To return to the command line, press the **q** key.

# Adding Comments

You can precede a line with a colon (:) to create a comment. However, the comment only appears in the command history buffer and not in the configuration. Therefore, you can view the comment with the **show history** command or by pressing an arrow key to retrieve a previous command, but because the comment is not in the configuration, the **write terminal** command does not display it.

## Text Configuration Files

This section describes how to format a text configuration file that you can download to the security appliance, and includes the following topics:

- [How Commands Correspond with Lines in the Text File, page C-6](#)
- [Command-Specific Configuration Mode Commands, page C-6](#)
- [Automatic Text Entries, page C-7](#)
- [Line Order, page C-7](#)
- [Commands Not Included in the Text Configuration, page C-7](#)
- [Passwords, page C-7](#)
- [Multiple Security Context Files, page C-7](#)

## How Commands Correspond with Lines in the Text File

The text configuration file includes lines that correspond with the commands described in this guide.

In examples, commands are preceded by a CLI prompt. The prompt in the following example is “hostname(config)#”:

```
hostname(config)# context a
```

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

```
context a
```

## Command-Specific Configuration Mode Commands

Command-specific configuration mode commands appear indented under the main command when entered at the command line. Your text file lines do not need to be indented, as long as the commands appear directly following the main command. For example, the following unindented text is read the same as indented text:

```
interface gigabitethernet0/0
nameif inside
interface gigabitethernet0/1
 nameif outside
```



## Automatic Text Entries

When you download a configuration to the security appliance, the security appliance inserts some lines automatically. For example, the security appliance inserts lines for default settings or for the time the configuration was modified. You do not need to enter these automatic entries when you create your text file.

## Line Order

For the most part, commands can be in any order in the file. However, some lines, such as ACEs, are processed in the order they appear, and the order can affect the function of the access list. Other commands might also have order requirements. For example, you must enter the **nameif** command for an interface first because many subsequent commands use the name of the interface. Also, commands in a command-specific configuration mode must directly follow the main command.

## Commands Not Included in the Text Configuration

Some commands do not insert lines in the configuration. For example, a runtime command such as **show running-config** does not have a corresponding line in the text file.

## Passwords

The login, enable, and user passwords are automatically encrypted before they are stored in the configuration. For example, the encrypted form of the password “cisco” might look like jMorNbK0514fadBh. You can copy the configuration passwords to another security appliance in their encrypted form, but you cannot unencrypt the passwords yourself.

If you enter an unencrypted password in a text file, the security appliance does not automatically encrypt them when you copy the configuration to the security appliance. The security appliance only encrypts them when you save the running configuration from the command line using the **copy running-config startup-config** or **write memory** command.

## Multiple Security Context Files

For multiple security contexts, the entire configuration consists of multiple parts:

- The security context configurations
- The system configuration, which identifies basic settings for the security appliance, including a list of contexts
- The admin context, which provides network interfaces for the system configuration

The system configuration does not include any interfaces or network settings for itself. Rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses a context that is designated as the admin context.

Each context is similar to a single context mode configuration. The system configuration differs from a context configuration in that the system configuration includes system-only commands (such as a list of all contexts) while other typical commands are not present (such as many interface parameters).





# APPENDIX **D**

## Addresses, Protocols, and Ports

---

This appendix provides a quick reference for IP addresses, protocols, and applications. This appendix includes the following sections:

- [IPv4 Addresses and Subnet Masks, page D-1](#)
- [IPv6 Addresses, page D-5](#)
- [Protocols and Applications, page D-11](#)
- [TCP and UDP Ports, page D-11](#)
- [Local Ports and Protocols, page D-14](#)
- [ICMP Types, page D-15](#)

## IPv4 Addresses and Subnet Masks

This section describes how to use IPv4 addresses in the security appliance. An IPv4 address is a 32-bit number written in dotted-decimal notation: four 8-bit fields (octets) converted from binary to decimal numbers, separated by dots. The first part of an IP address identifies the network on which the host resides, while the second part identifies the particular host on the given network. The network number field is called the network prefix. All hosts on a given network share the same network prefix but must have a unique host number. In classful IP, the class of the address determines the boundary between the network prefix and the host number.

This section includes the following topics:

- [Classes, page D-1](#)
- [Private Networks, page D-2](#)
- [Subnet Masks, page D-2](#)

## Classes

IP host addresses are divided into three different address classes: Class A, Class B, and Class C. Each class fixes the boundary between the network prefix and the host number at a different point within the 32-bit address. Class D addresses are reserved for multicast IP.

- Class A addresses (1.xxx.xxx.xxx through 126.xxx.xxx.xxx) use only the first octet as the network prefix.

- Class B addresses (128.0.xxx.xxx through 191.255.xxx.xxx) use the first two octets as the network prefix.
- Class C addresses (192.0.0.xxx through 223.255.255.xxx) use the first three octets as the network prefix.

Because Class A addresses have 16,777,214 host addresses, and Class B addresses 65,534 hosts, you can use subnet masking to break these huge networks into smaller subnets.

## Private Networks

If you need large numbers of addresses on your network, and they do not need to be routed on the Internet, you can use private IP addresses that the Internet Assigned Numbers Authority (IANA) recommends (see RFC 1918). The following address ranges are designated as private networks that should not be advertised:

- 10.0.0.0 through 10.255.255.255
- 172.16.0.0 through 172.31.255.255
- 192.168.0.0 through 192.168.255.255

## Subnet Masks

A subnet mask lets you convert a single Class A, B, or C network into multiple networks. With a subnet mask, you can create an extended network prefix that adds bits from the host number to the network prefix. For example, a Class C network prefix always consists of the first three octets of the IP address. But a Class C extended network prefix uses part of the fourth octet as well.

Subnet masking is easy to understand if you use binary notation instead of dotted decimal. The bits in the subnet mask have a one-to-one correspondence with the Internet address:

- The bits are set to 1 if the corresponding bit in the IP address is part of the extended network prefix.
- The bits are set to 0 if the bit is part of the host number.

**Example 1:** If you have the Class B address 129.10.0.0 and you want to use the entire third octet as part of the extended network prefix instead of the host number, you must specify a subnet mask of 11111111.11111111.11111111.00000000. This subnet mask converts the Class B address into the equivalent of a Class C address, where the host number consists of the last octet only.

**Example 2:** If you want to use only part of the third octet for the extended network prefix, then you must specify a subnet mask like 11111111.11111111.11111000.00000000, which uses only 5 bits of the third octet for the extended network prefix.

You can write a subnet mask as a dotted-decimal mask or as a */bits* (“slash *bits*”) mask. In Example 1, for a dotted-decimal mask, you convert each binary octet into a decimal number: 255.255.255.0. For a */bits* mask, you add the number of 1s: /24. In Example 2, the decimal number is 255.255.248.0 and the */bits* is /21.

You can also supernet multiple Class C networks into a larger network by using part of the third octet for the extended network prefix. For example, 192.168.0.0/20.

This section includes the following topics:

- [Determining the Subnet Mask, page D-3](#)
- [Determining the Address to Use with the Subnet Mask, page D-3](#)

## Determining the Subnet Mask

To determine the subnet mask based on how many hosts you want, see [Table D-1](#).

**Table D-1** *Hosts, Bits, and Dotted-Decimal Masks*

| Hosts <sup>1</sup> | /Bits Mask | Dotted-Decimal Mask                 |
|--------------------|------------|-------------------------------------|
| 16,777,216         | /8         | 255.0.0.0 Class A Network           |
| 65,536             | /16        | 255.255.0.0 Class B Network         |
| 32,768             | /17        | 255.255.128.0                       |
| 16,384             | /18        | 255.255.192.0                       |
| 8192               | /19        | 255.255.224.0                       |
| 4096               | /20        | 255.255.240.0                       |
| 2048               | /21        | 255.255.248.0                       |
| 1024               | /22        | 255.255.252.0                       |
| 512                | /23        | 255.255.254.0                       |
| 256                | /24        | 255.255.255.0 Class C Network       |
| 128                | /25        | 255.255.255.128                     |
| 64                 | /26        | 255.255.255.192                     |
| 32                 | /27        | 255.255.255.224                     |
| 16                 | /28        | 255.255.255.240                     |
| 8                  | /29        | 255.255.255.248                     |
| 4                  | /30        | 255.255.255.252                     |
| Do not use         | /31        | 255.255.255.254                     |
| 1                  | /32        | 255.255.255.255 Single Host Address |

1. The first and last number of a subnet are reserved, except for /32, which identifies a single host.

## Determining the Address to Use with the Subnet Mask

The following sections describe how to determine the network address to use with a subnet mask for a Class C-size and a Class B-size network. This section includes the following topics:

- [Class C-Size Network Address, page D-3](#)
- [Class B-Size Network Address, page D-4](#)

### Class C-Size Network Address

For a network between 2 and 254 hosts, the fourth octet falls on a multiple of the number of host addresses, starting with 0. For example, the 8-host subnets (/29) of 192.168.0.x are as follows:

| Subnet with Mask /29 (255.255.255.248) | Address Range <sup>1</sup>  |
|----------------------------------------|-----------------------------|
| 192.168.0.0                            | 192.168.0.0 to 192.168.0.7  |
| 192.168.0.8                            | 192.168.0.8 to 192.168.0.15 |

| Subnet with Mask /29 (255.255.255.248) | Address Range <sup>1</sup>     |
|----------------------------------------|--------------------------------|
| 192.168.0.16                           | 192.168.0.16 to 192.168.0.31   |
| ...                                    | ...                            |
| 192.168.0.248                          | 192.168.0.248 to 192.168.0.255 |

1. The first and last address of a subnet are reserved. In the first subnet example, you cannot use 192.168.0.0 or 192.168.0.7.

## Class B-Size Network Address

To determine the network address to use with the subnet mask for a network with between 254 and 65,534 hosts, you need to determine the value of the third octet for each possible extended network prefix. For example, you might want to subnet an address like 10.1.x.0, where the first two octets are fixed because they are used in the extended network prefix, and the fourth octet is 0 because all bits are used for the host number.

To determine the value of the third octet, follow these steps:

- Step 1** Calculate how many subnets you can make from the network by dividing 65,536 (the total number of addresses using the third and fourth octet) by the number of host addresses you want.
- For example, 65,536 divided by 4096 hosts equals 16.
- Therefore, there are 16 subnets of 4096 addresses each in a Class B-size network.
- Step 2** Determine the multiple of the third octet value by dividing 256 (the number of values for the third octet) by the number of subnets:
- In this example,  $256/16 = 16$ .
- The third octet falls on a multiple of 16, starting with 0.
- Therefore, the 16 subnets of the network 10.1 are as follows:

| Subnet with Mask /20 (255.255.240.0) | Address Range <sup>1</sup> |
|--------------------------------------|----------------------------|
| 10.1.0.0                             | 10.1.0.0 to 10.1.15.255    |
| 10.1.16.0                            | 10.1.16.0 to 10.1.31.255   |
| 10.1.32.0                            | 10.1.32.0 to 10.1.47.255   |
| ...                                  | ...                        |
| 10.1.240.0                           | 10.1.240.0 to 10.1.255.255 |

1. The first and last address of a subnet are reserved. In the first subnet example, you cannot use 10.1.0.0 or 10.1.15.255.

# IPv6 Addresses

IPv6 is the next generation of the Internet Protocol after IPv4. It provides an expanded address space, a simplified header format, improved support for extensions and options, flow labeling capability, and authentication and privacy capabilities. IPv6 is described in RFC 2460. The IPv6 addressing architecture is described in RFC 3513.

This section describes the IPv6 address format and architecture and includes the following topics:

- [IPv6 Address Format, page D-5](#)
- [IPv6 Address Types, page D-6](#)
- [IPv6 Address Prefixes, page D-10](#)

**Note**

This section describes the IPv6 address format, the types, and prefixes. For information about configuring the security appliance to use IPv6, see [Chapter 7, “Configuring Interface Parameters.”](#)

## IPv6 Address Format

IPv6 addresses are represented as a series of eight 16-bit hexadecimal fields separated by colons (:) in the format: x:x:x:x:x:x:x:x. The following are two examples of IPv6 addresses:

- 2001:0DB8:7654:3210:FEDC:BA98:7654:3210
- 2001:0DB8:0000:0000:0008:0800:200C:417A

**Note**

The hexadecimal letters in IPv6 addresses are not case-sensitive.

It is not necessary to include the leading zeros in an individual field of the address. But each field must contain at least one digit. So the example address 2001:0DB8:0000:0000:0008:0800:200C:417A can be shortened to 2001:0DB8:0:0:8:800:200C:417A by removing the leading zeros from the third through sixth fields from the left. The fields that contained all zeros (the third and fourth fields from the left) were shortened to a single zero. The fifth field from the left had the three leading zeros removed, leaving a single 8 in that field, and the sixth field from the left had the one leading zero removed, leaving 800 in that field.

It is common for IPv6 addresses to contain several consecutive hexadecimal fields of zeros. You can use two colons (::) to compress consecutive fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent the successive hexadecimal fields of zeros). [Table D-2](#) shows several examples of address compression for different types of IPv6 address.

**Table D-2** IPv6 Address Compression Examples

| Address Type | Standard Form               | Compressed Form        |
|--------------|-----------------------------|------------------------|
| Unicast      | 2001:0DB8:0:0:0:BA98:0:3210 | 2001:0DB8::BA98:0:3210 |
| Multicast    | FF01:0:0:0:0:0:0:101        | FF01::101              |
| Loopback     | 0:0:0:0:0:0:0:1             | ::1                    |
| Unspecified  | 0:0:0:0:0:0:0:0             | ::                     |

**Note**

Two colons (::) can be used only once in an IPv6 address to represent successive fields of zeros.

An alternative form of the IPv6 format is often used when dealing with an environment that contains both IPv4 and IPv6 addresses. This alternative has the format `x:x:x:x:x:y.y.y.y`, where `x` represent the hexadecimal values for the six high-order parts of the IPv6 address and `y` represent decimal values for the 32-bit IPv4 part of the address (which takes the place of the remaining two 16-bit parts of the IPv6 address). For example, the IPv4 address 192.168.1.1 could be represented as the IPv6 address `0:0:0:0:0:FFFF:192.168.1.1`, or `::FFFF:192.168.1.1`.

## IPv6 Address Types

The following are the three main types of IPv6 addresses:

- **Unicast**—A unicast address is an identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address. An interface may have more than one unicast address assigned to it.
- **Multicast**—A multicast address is an identifier for a set of interfaces. A packet sent to a multicast address is delivered to all addresses identified by that address.
- **Anycast**—An anycast address is an identifier for a set of interfaces. Unlike a multicast address, a packet sent to an anycast address is only delivered to the “nearest” interface, as determined by the measure of distances for the routing protocol.

**Note**

There are no broadcast addresses in IPv6. Multicast addresses provide the broadcast functionality.

This section includes the following topics:

- [Unicast Addresses, page D-6](#)
- [Multicast Address, page D-8](#)
- [Anycast Address, page D-9](#)
- [Required Addresses, page D-10](#)

## Unicast Addresses

This section describes IPv6 unicast addresses. Unicast addresses identify an interface on a network node.

This section includes the following topics:

- [Global Address, page D-7](#)
- [Site-Local Address, page D-7](#)
- [Link-Local Address, page D-7](#)
- [IPv4-Compatible IPv6 Addresses, page D-7](#)
- [Unspecified Address, page D-8](#)
- [Loopback Address, page D-8](#)
- [Interface Identifiers, page D-8](#)



## Global Address

The general format of an IPv6 global unicast address is a global routing prefix followed by a subnet ID followed by an interface ID. The global routing prefix can be any prefix not reserved by another IPv6 address type (see [IPv6 Address Prefixes, page D-10](#), for information about the IPv6 address type prefixes).

All global unicast addresses, other than those that start with binary 000, have a 64-bit interface ID in the Modified EUI-64 format. See [Interface Identifiers, page D-8](#), for more information about the Modified EUI-64 format for interface identifiers.

Global unicast address that start with the binary 000 do not have any constraints on the size or structure of the interface ID portion of the address. One example of this type of address is an IPv6 address with an embedded IPv4 address (see [IPv4-Compatible IPv6 Addresses, page D-7](#)).

## Site-Local Address

Site-local addresses are used for addressing within a site. They can be use to address an entire site without using a globally unique prefix. Site-local addresses have the prefix FEC0::/10, followed by a 54-bit subnet ID, and end with a 64-bit interface ID in the modified EUI-64 format.

Site-local Routers do not forward any packets that have a site-local address for a source or destination outside of the site. Therefore, site-local addresses can be considered private addresses.

## Link-Local Address

All interfaces are required to have at least one link-local address. You can configure multiple IPv6 addresses per interfaces, but only one link-local address.

A link-local address is an IPv6 unicast address that can be automatically configured on any interface using the link-local prefix FE80::/10 and the interface identifier in modified EUI-64 format. Link-local addresses are used in the neighbor discovery protocol and the stateless autoconfiguration process. Nodes with a link-local address can communicate; they do not need a site-local or globally unique address to communicate.

Routers do not forward any packets that have a link-local address for a source or destination. Therefore, link-local addresses can be considered private addresses.

## IPv4-Compatible IPv6 Addresses

There are two types of IPv6 addresses that can contain IPv4 addresses.

The first type is the “IPv4-compatibly IPv6 address.” The IPv6 transition mechanisms include a technique for hosts and routers to dynamically tunnel IPv6 packets over IPv4 routing infrastructure. IPv6 nodes that use this technique are assigned special IPv6 unicast addresses that carry a global IPv4 address in the low-order 32 bits. This type of address is termed an “IPv4-compatible IPv6 address” and has the format ::y.y.y.y, where y.y.y.y is an IPv4 unicast address.



### Note

The IPv4 address used in the “IPv4-compatible IPv6 address” must be a globally-unique IPv4 unicast address.

The second type of IPv6 address which holds an embedded IPv4 address is called the “IPv4-mapped IPv6 address.” This address type is used to represent the addresses of IPv4 nodes as IPv6 addresses. This type of address has the format ::FFFF:y.y.y.y, where y.y.y.y is an IPv4 unicast address.

## Unspecified Address

The unspecified address, 0:0:0:0:0:0:0:0, indicates the absence of an IPv6 address. For example, a newly initialized node on an IPv6 network may use the unspecified address as the source address in its packets until it receives its IPv6 address.

**Note**

The IPv6 unspecified address cannot be assigned to an interface. The unspecified IPv6 addresses must not be used as destination addresses in IPv6 packets or the IPv6 routing header.

## Loopback Address

The loopback address, 0:0:0:0:0:0:0:1, may be used by a node to send an IPv6 packet to itself. The loopback address in IPv6 functions the same as the loopback address in IPv4 (127.0.0.1).

**Note**

The IPv6 loopback address cannot be assigned to a physical interface. A packet that has the IPv6 loopback address as its source or destination address must remain within the node that created the packet. IPv6 routers do not forward packets that have the IPv6 loopback address as their source or destination address.

## Interface Identifiers

Interface identifiers in IPv6 unicast addresses are used to identify the interfaces on a link. They need to be unique within a subnet prefix. In many cases, the interface identifier is derived from the interface link-layer address. The same interface identifier may be used on multiple interfaces of a single node, as long as those interfaces are attached to different subnets.


For all unicast addresses, except those that start with the binary 000, the interface identifier is required to be 64 bits long and to be constructed in the Modified EUI-64 format. The Modified EUI-64 format is created from the 48-bit MAC address by inverting the universal/local bit in the address and by inserting the hexadecimal number FFFE between the upper three bytes and lower three bytes of the of the MAC address.

For example, an interface with the MAC address of 00E0.b601.3B7A would have a 64-bit interface ID of 02E0:B6FF:FE01:3B7A.

## Multicast Address

An IPv6 multicast address is an identifier for a group of interfaces, typically on different nodes. A packet sent to a multicast address is delivered to all interfaces identified by the multicast address. An interface may belong to any number of multicast groups.

An IPv6 multicast address has a prefix of FF00::/8 (1111 1111). The octet following the prefix defines the type and scope of the multicast address. A permanently assigned (“well known”) multicast address has a flag parameter equal to 0; a temporary (“transient”) multicast address has a flag parameter equal to 1. A multicast address that has the scope of a node, link, site, or organization, or a global scope has a scope parameter of 1, 2, 5, 8, or E, respectively. For example, a multicast address with the prefix FF02::/16 is a permanent multicast address with a link scope. [Figure D-1](#) shows the format of the IPv6 multicast address.

**Figure D-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 security appliance.

## Required Addresses

IPv6 hosts must, at a minimum, be configured with the following addresses (either automatically or manually):

- A link-local address for each interface.
- The loopback address.
- The All-Nodes multicast addresses
- A Solicited-Node multicast address for each unicast or anycast address.

IPv6 routers must, at a minimum, be configured with the following addresses (either automatically or manually):

- The required host addresses.
- The Subnet-Router anycast addresses for all interfaces for which it is configured to act as a router.
- The All-Routers multicast addresses.

## IPv6 Address Prefixes

An IPv6 address prefix, in the format `ipv6-prefix/prefix-length`, can be used to represent bit-wise contiguous blocks of the entire address space. The IPv6-prefix must be in the form documented in RFC 2373 where the address is specified in hexadecimal using 16-bit values between colons. The prefix length is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). For example, `2001:0DB8:8086:6502::/32` is a valid IPv6 prefix.

The IPv6 prefix identifies the type of IPv6 address. [Table D-3](#) shows the prefixes for each IPv6 address type.

**Table D-3**      *IPv6 Address Type Prefixes*

| Address Type         | Binary Prefix                         | IPv6 Notation |
|----------------------|---------------------------------------|---------------|
| Unspecified          | 000...0 (128 bits)                    | ::/128        |
| Loopback             | 000...1 (128 bits)                    | ::1/128       |
| Multicast            | 11111111                              | FF00::/8      |
| Link-Local (unicast) | 1111111010                            | FE80::/10     |
| Site-Local (unicast) | 1111111111                            | FEC0::/10     |
| Global (unicast)     | All other addresses.                  |               |
| Anycast              | Taken from the unicast address space. |               |

# Protocols and Applications

Table D-4 lists the protocol literal values and port numbers; either can be entered in security appliance commands.

**Table D-4 Protocol Literal Values**

| Literal | Value | Description                                                                                                               |
|---------|-------|---------------------------------------------------------------------------------------------------------------------------|
| ah      | 51    | Authentication Header for IPv6, RFC 1826.                                                                                 |
| eigrp   | 88    | Enhanced Interior Gateway Routing Protocol.                                                                               |
| esp     | 50    | Encapsulated Security Payload for IPv6, RFC 1827.                                                                         |
| gre     | 47    | Generic Routing Encapsulation.                                                                                            |
| icmp    | 1     | Internet Control Message Protocol, RFC 792.                                                                               |
| icmp6   | 58    | Internet Control Message Protocol for IPv6, RFC 2463.                                                                     |
| igmp    | 2     | Internet Group Management Protocol, RFC 1112.                                                                             |
| igrp    | 9     | Interior Gateway Routing Protocol.                                                                                        |
| ip      | 0     | Internet Protocol.                                                                                                        |
| ipinip  | 4     | IP-in-IP encapsulation.                                                                                                   |
| ipsec   | 50    | IP Security. Entering the ipsec protocol literal is equivalent to entering the esp protocol literal.                      |
| nos     | 94    | Network Operating System (Novell's NetWare).                                                                              |
| ospf    | 89    | Open Shortest Path First routing protocol, RFC 1247.                                                                      |
| pcp     | 108   | Payload Compression Protocol.                                                                                             |
| pim     | 103   | Protocol Independent Multicast.                                                                                           |
| pptp    | 47    | Point-to-Point Tunneling Protocol. Entering the pptp protocol literal is equivalent to entering the gre protocol literal. |
| snp     | 109   | Sitara Networks Protocol.                                                                                                 |
| tcp     | 6     | Transmission Control Protocol, RFC 793.                                                                                   |
| udp     | 17    | User Datagram Protocol, RFC 768.                                                                                          |

Protocol numbers can be viewed online at the IANA website:

<http://www.iana.org/assignments/protocol-numbers>

## TCP and UDP Ports

Table D-5 lists the literal values and port numbers; either can be entered in security appliance commands. See the following caveats:

- The security appliance uses port 1521 for SQL\*Net. This is the default port used by Oracle for SQL\*Net. This value, however, does not agree with IANA port assignments.

- The security appliance listens for RADIUS on ports 1645 and 1646. If your RADIUS server uses the standard ports 1812 and 1813, you can configure the security appliance to listen to those ports using the **authentication-port** and **accounting-port** commands.
- To assign a port for DNS access, use the **domain** literal value, not **dns**. If you use **dns**, the security appliance assumes you meant to use the **dnsix** literal value.

Port numbers can be viewed online at the IANA website:

<http://www.iana.org/assignments/port-numbers>

**Table D-5** Port Literal Values

| Literal    | TCP or UDP? | Value | Description                                                                |
|------------|-------------|-------|----------------------------------------------------------------------------|
| aol        | TCP         | 5190  | America Online                                                             |
| bgp        | TCP         | 179   | Border Gateway Protocol, RFC 1163                                          |
| biff       | UDP         | 512   | Used by mail system to notify users that new mail is received              |
| bootpc     | UDP         | 68    | Bootstrap Protocol Client                                                  |
| bootps     | UDP         | 67    | Bootstrap Protocol Server                                                  |
| chargen    | TCP         | 19    | Character Generator                                                        |
| citrix-ica | TCP         | 1494  | Citrix Independent Computing Architecture (ICA) protocol                   |
| cmd        | TCP         | 514   | Similar to <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 signalling                                                      |
| hostname   | TCP         | 101   | NIC Host Name Server                                                       |
| ident      | TCP         | 113   | Ident authentication service                                               |
| imap4      | TCP         | 143   | Internet Message Access Protocol, version 4                                |
| irc        | TCP         | 194   | Internet Relay Chat protocol                                               |

**Table D-5** Port Literal Values (continued)

| Literal           | TCP or UDP? | Value | Description                                                       |
|-------------------|-------------|-------|-------------------------------------------------------------------|
| isakmp            | UDP         | 500   | Internet Security Association and Key Management Protocol         |
| kerberos          | TCP, UDP    | 750   | Kerberos                                                          |
| klogin            | TCP         | 543   | KLOGIN                                                            |
| kshell            | TCP         | 544   | Korn Shell                                                        |
| ldap              | TCP         | 389   | Lightweight Directory Access Protocol                             |
| ldaps             | TCP         | 636   | Lightweight Directory Access Protocol (SSL)                       |
| lpd               | TCP         | 515   | Line Printer Daemon - printer spooler                             |
| login             | TCP         | 513   | Remote login                                                      |
| lotusnotes        | TCP         | 1352  | IBM Lotus Notes                                                   |
| mobile-ip         | UDP         | 434   | MobileIP-Agent                                                    |
| nameserver        | UDP         | 42    | Host Name Server                                                  |
| netbios-ns        | UDP         | 137   | NetBIOS Name Service                                              |
| netbios-dgm       | UDP         | 138   | NetBIOS Datagram Service                                          |
| netbios-ssn       | TCP         | 139   | NetBIOS Session Service                                           |
| nntp              | TCP         | 119   | Network News Transfer Protocol                                    |
| ntp               | UDP         | 123   | Network Time Protocol                                             |
| pcanywhere-status | UDP         | 5632  | pcAnywhere status                                                 |
| pcanywhere-data   | TCP         | 5631  | pcAnywhere data                                                   |
| pim-auto-rp       | TCP, UDP    | 496   | Protocol Independent Multicast, reverse path flooding, dense mode |
| pop2              | TCP         | 109   | Post Office Protocol - Version 2                                  |
| pop3              | TCP         | 110   | Post Office Protocol - Version 3                                  |
| pptp              | TCP         | 1723  | Point-to-Point Tunneling Protocol                                 |
| radius            | UDP         | 1645  | Remote Authentication Dial-In User Service                        |
| radius-acct       | UDP         | 1646  | Remote Authentication Dial-In User Service (accounting)           |
| rip               | UDP         | 520   | Routing Information Protocol                                      |
| secureid-udp      | UDP         | 5510  | SecureID over UDP                                                 |
| smtp              | TCP         | 25    | Simple Mail Transport Protocol                                    |
| snmp              | UDP         | 161   | Simple Network Management Protocol                                |
| snmptrap          | UDP         | 162   | Simple Network Management Protocol - Trap                         |
| sqlnet            | TCP         | 1521  | Structured Query Language Network                                 |
| ssh               | TCP         | 22    | Secure Shell                                                      |
| sunrpc (rpc)      | TCP, UDP    | 111   | Sun Remote Procedure Call                                         |
| syslog            | UDP         | 514   | System Log                                                        |

**Table D-5** Port Literal Values (continued)

| Literal | TCP or UDP? | Value | Description                                           |
|---------|-------------|-------|-------------------------------------------------------|
| tacacs  | TCP, UDP    | 49    | Terminal Access Controller Access Control System Plus |
| talk    | TCP, UDP    | 517   | Talk                                                  |
| telnet  | TCP         | 23    | RFC 854 Telnet                                        |
| tftp    | UDP         | 69    | Trivial File Transfer Protocol                        |
| time    | UDP         | 37    | Time                                                  |
| uucp    | TCP         | 540   | UNIX-to-UNIX Copy Program                             |
| who     | UDP         | 513   | Who                                                   |
| whois   | TCP         | 43    | Who Is                                                |
| www     | TCP         | 80    | World Wide Web                                        |
| xdmcp   | UDP         | 177   | X Display Manager Control Protocol                    |

## Local Ports and Protocols

Table D-6 lists the protocols, TCP ports, and UDP ports that the security appliance may open to process traffic destined to the security appliance. Unless you enable the features and services listed in Table D-6, the security appliance does *not* open any local protocols or any TCP or UDP ports. You must configure a feature or service for the security appliance to open the default listening protocol or port. In many cases you can configure ports other than the default port when you enable a feature or service.

**Table D-6** Protocols and Ports Opened by Features and Services

| Feature or Service                                | Protocol | Port Number | Comments                                                                                   |
|---------------------------------------------------|----------|-------------|--------------------------------------------------------------------------------------------|
| DHCP                                              | UDP      | 67,68       | —                                                                                          |
| Failover Control                                  | 108      | N/A         | —                                                                                          |
| HTTP                                              | TCP      | 80          | —                                                                                          |
| HTTPS                                             | TCP      | 443         | —                                                                                          |
| ICMP                                              | 1        | N/A         | —                                                                                          |
| IGMP                                              | 2        | N/A         | Protocol only open on destination IP address 224.0.0.1                                     |
| ISAKMP/IKE                                        | UDP      | 500         | Configurable.                                                                              |
| IPSec (ESP)                                       | 50       | N/A         | —                                                                                          |
| IPSec over UDP (NAT-T)                            | UDP      | 4500        | —                                                                                          |
| IPSec over UDP (Cisco VPN 3000 Series compatible) | UDP      | 10000       | Configurable.                                                                              |
| IPSec over TCP (CTCP)                             | TCP      | —           | No default port is used. You must specify the port number when configuring IPSec over TCP. |



**Table D-6** *Protocols and Ports Opened by Features and Services (continued)*

| Feature or Service                       | Protocol | Port Number | Comments                                                             |
|------------------------------------------|----------|-------------|----------------------------------------------------------------------|
| NTP                                      | UDP      | 123         | —                                                                    |
| OSPF                                     | 89       | N/A         | Protocol only open on destination IP address 224.0.0.5 and 224.0.0.6 |
| PIM                                      | 103      | N/A         | Protocol only open on destination IP address 224.0.0.13              |
| RIP                                      | UDP      | 520         | —                                                                    |
| RIPv2                                    | UDP      | 520         | Port only open on destination IP address 224.0.0.9                   |
| SNMP                                     | UDP      | 161         | Configurable.                                                        |
| SSH                                      | TCP      | 22          | —                                                                    |
| Stateful Update                          | 105      | N/A         | —                                                                    |
| Telnet                                   | TCP      | 23          | —                                                                    |
| VPN Load Balancing                       | UDP      | 9023        | Configurable.                                                        |
| VPN Individual User Authentication Proxy | UDP      | 1645, 1646  | Port accessible only over VPN tunnel.                                |

## ICMP Types

Table D-7 lists the ICMP type numbers and names that you can enter in security appliance commands:

**Table D-7** *ICMP Types*

| ICMP Number | ICMP Name            |
|-------------|----------------------|
| 0           | echo-reply           |
| 3           | unreachable          |
| 4           | source-quench        |
| 5           | redirect             |
| 6           | alternate-address    |
| 8           | echo                 |
| 9           | router-advertisement |
| 10          | router-solicitation  |
| 11          | time-exceeded        |
| 12          | parameter-problem    |
| 13          | timestamp-request    |
| 14          | timestamp-reply      |
| 15          | information-request  |
| 16          | information-reply    |
| 17          | mask-request         |

**Table D-7** *ICMP Types (continued)*

| ICMP Number | ICMP Name        |
|-------------|------------------|
| 18          | mask-reply       |
| 31          | conversion-error |
| 32          | mobile-redirect  |



## APPENDIX **E**

# Configuring an External Server for Authorization and Authentication

---

This appendix describes how to configure an external LDAP or RADIUS server to support the authentication and authorization of security appliance, VPN3000, and PIX users. Authentication determines who the user is and authorization determines what the user can do. Before you configure the security appliance to use an external server, you must configure the server with the correct security appliance authorization attributes and, from a subset of these attributes, assign specific permissions to individual users.

This appendix includes the following sections:

- [Selecting LDAP, RADIUS, or Local Authentication and Authorization](#)
- [Understanding Policy Enforcement of Permissions and Attributes](#)
- [Configuring an External LDAP Server](#)
- [Configuring an External RADIUS Server](#)

## Selecting LDAP, RADIUS, or Local Authentication and Authorization

To help you decide which authentication or authorization method is right for your platform, this section describes the LDAP and RADIUS support provided with the security appliance (ASA), PIX, and the VPN 3000 platforms.

- **LDAP Authentication**  
Supported on PIX 7.1.x and the security appliance only. VPN 3000 does not support native LDAP authentication. The LDAP server retrieves and searches for the username and enforces any defined attributes as part of the authorization function.
- **LDAP Authorization**  
Supported on PIX, VPN 3000, and the security appliance. The LDAP server retrieves/searches the username and enforces any defined attributes.
- **RADIUS Authentication**  
Supported on PIX, VPN 3000, and the security appliance. The RADIUS server retrieves/searches the username and enforces any defined attributes as it performs the authorization function.
- **RADIUS Authorization**

Supported on PIX, VPN 3000, and the security appliance. The RADIUS server retrieves/searches the username and enforces any defined attributes.

- Local Authentication

Supported on PIX, VPN 3000, and the security appliance. The Local/Internal server retrieves/searches the username and enforces any defined attributes as part of the authorization function.

- Local Authorization

Supported on PIX 7.1.x and the security appliance only. The Local/Internal server retrieves/searches the username and enforces any defined attributes.

## Understanding Policy Enforcement of Permissions and Attributes

You can configure the security appliance to receive user attributes from either the LOCAL/internal database, a RADIUS/LDAP authentication server, or a RADIUS/LDAP authorization server. You can also place users into group-policies with different attributes, but the user attributes will always take precedence. After the device authenticates the user and group(s), the security appliance combines the user and group attribute sets into one aggregate attribute set. The security appliance uses the attributes in the following order and applies the aggregate attribute set to the authenticated user.

1. User attributes—The server returns these after successful user authentication or authorization. These take precedence over all others.
2. Group policy attributes—These attributes come from the group policy associated with the user. You identify the user group policy name in the local database by the 'vpn-group-policy' attribute or from an external RADIUS/LDAP server by the value of the RADIUS CLASS attribute (25) in the format 'OU=GroupName;'. The group policy provides any attributes that are missing from the user attributes. User attributes override group policy attributes if both have a value.
3. Tunnel group default-group-policy attributes—These attributes come from the default-group-policy (Base group) that is associated with the tunnel group. After a lookup of that group policy, the Tunnel Group's default-group-policy provide any attributes that are missing from the user or group policy attributes. User attributes override group policy attributes if both have a value.
4. System default attributes—System default attributes provide any attributes that are missing from the user, group, or tunnel group attributes.

## Configuring an External LDAP Server



### Note

For more information on the LDAP protocol, see RFCs 1777, 2251, and 2849.

This section describes the structure, schema, and attributes of an LDAP server. It includes the following topics:

- [Reviewing the LDAP Directory Structure and Configuration Procedure](#)
- [Organizing the Security Appliance LDAP Schema](#)
- [Defining the Security Appliance LDAP Schema](#)

- [Loading the Schema in the LDAP Server](#)
- [Defining User Permissions](#)
- [Reviewing Examples of Active Directory Configurations](#)

## Reviewing the LDAP Directory Structure and Configuration Procedure

An LDAP server stores information as entries in a directory. An LDAP schema defines what types of information such entries store. The schema lists classes and the set of (required and optional) attributes that objects of each class can contain.

To configure your LDAP server to interoperate with the security appliance, define a security appliance authorization schema. A security appliance authorization schema defines the class and attributes of that class that the security appliance supports. Specifically, it comprises the object class (cVPN3000-User-Authorization) and all its possible attributes that may be used to authorize a security appliance user (such as access hours, primary DNS, and so on). Each attribute comprises the attribute name, its number (called an object identifier or OID), its type, and its possible values.

Once you have defined the security appliance authorization schema and loaded it on your server, define the security appliance attributes and permissions and their respective values for each user who will be authorizing to the server.

In summary, to set up your LDAP server:

- Design your security appliance LDAP authorization schema based on the hierarchical set-up of your organization
- Define the security appliance authorization schema
- Load the schema on the LDAP server
- Define permissions for each user on the LDAP server


The specific steps of these processes vary, depending on which type of LDAP server you are using.

## Organizing the Security Appliance LDAP Schema

Before you actually create your schema, think about how your organization is structured. Your LDAP schema should reflect the logical hierarchy of your organization.

For example, suppose an employee at your company, Example Corporation, is named Terry. Terry works in the Engineering group. Your LDAP hierarchy could have one or many levels. You might decide to set up a shallow, single-level hierarchy in which Terry is considered a member of Example Corporation. Or, you could set up a multi-level hierarchy in which Terry is considered to be a member of the department Engineering, which is a member of an organizational unit called People, which is itself a member of Example Corporation. See [Figure E-1](#) for an example of this multi-level hierarchy.

A multi-level hierarchy has more granularity, but a single level hierarchy is quicker to search.

**Figure E-1 A Multi-Level LDAP Hierarchy**

## Searching the Hierarchy

The security appliance lets you tailor the search within the LDAP hierarchy. You configure the following three fields on the security appliance to define where in the LDAP hierarchy your search begins, its extent, and the type of information it is looking for. Together these fields allow you to limit the search of the hierarchy to just the part of the tree that contains the user permissions.

- **LDAP Base DN** defines where in the LDAP hierarchy the server should begin searching for user information when it receives an authorization request from the security appliance.
- **Search Scope** defines the extent of the search in the LDAP hierarchy. The search proceeds this many levels in the hierarchy below the LDAP Base DN. You can choose to have the server search only the level immediately below, or it can search the entire subtree. A single level search is quicker, but a subtree search is more extensive.
- **Naming Attribute(s)** defines the Relative Distinguished Name (RDN) that uniquely identifies an entry in the LDAP server. Common naming attributes are: cn (Common Name) and ui (user identification).

Figure E-1 shows a possible LDAP hierarchy for Example Corporation. Given this hierarchy, you could define your search in different ways. Table E-1 shows two possible search configurations.

In the first example configuration, when Terry establishes his or her IPsec tunnel with LDAP authorization required, the security appliance sends a search request to the LDAP server indicating it should search for Terry in the Engineering group. This search is quick.

In the second example configuration, the security appliance sends a search request indicating the server should search for Terry within Example Corporation. This search takes longer.

**Table E-1 Example Search Configurations**

| # | LDAP Base DN                                              | Search Scope | Naming Attribute | Result         |
|---|-----------------------------------------------------------|--------------|------------------|----------------|
| 1 | group= Engineering,ou=People,dc=ExampleCorporation,dc=com | One Level    | cn=Terry         | Quicker search |
| 2 | dc=ExampleCorporation,dc=com                              | Subtree      | cn=Terry         | Longer search  |

## Binding the Security Appliance to the LDAP Server

Some LDAP servers (including the Microsoft Active Directory server) require the security appliance to establish a handshake via authenticated binding before they accept requests for any other LDAP operations. The security appliance identifies itself for authenticated binding by attaching a Login DN field to the user authentication request. The Login DN field defines the authentication characteristics of the security appliance; these characteristics should correspond to those of a user with administration privileges. An example Login DN field could be: `cn=Administrator, cn=users, ou=people, dc=example, dc=com`.

## Defining the Security Appliance LDAP Schema

Once you have decided how to structure your user information in the LDAP hierarchy, define this organization in a schema. To define the schema, begin by defining the object class name. The class name for the security appliance directory is `cVPN3000-User-Authorization`. The class has the object identifier (OID) `1.2.840.113556.1.8000.795.1.1`. Every entry or user in the directory is an object of this class.

Some LDAP servers (for example, the Microsoft Active Directory LDAP server) do not allow you to reuse the class OID once you have defined it. Use the next incremental OID. For example, if you incorrectly defined the class name as `cVPN3000-Usr-Authorization` with OID `1.2.840.113556.1.8000.795.1.1`, you can enter the correct class name `cVPN3000-User-Authorization` with the next OID, for example, `1.2.840.113556.1.8000.795.1.2`.

For the Microsoft Active Directory LDAP server, define the schema in text form in a file using the LDAP Data Interchange Format (LDIF). This file has an extension of `.ldif`, for example: `schema.ldif`. Other LDAP servers use graphical user interfaces or script files to define the object class and its attributes. For more information on LDIF, see RFC-2849.



### Note

- All LDAP attributes for all three appliances begin with the letters `cVPN3000`; for example: `cVPN3000-Access-Hours`.
- The appliances enforce the LDAP attributes based on attribute name, not numeric ID. RADIUS attributes, on the other hand, are enforced by numeric ID, not by name.
- Authorization refers to the process of enforcing permissions or attributes. An LDAP server defined as an authentication or authorization server will enforce permissions or attributes if they are configured.

For a complete list of attributes for the security appliance, the PIX Firewall and the VPN 3000, see [Table E-2](#).

All strings are case-sensitive and you must use an attribute name as capitalized in the table even if it conflicts with how a term is typically written. For example, use `cVPN3000-IETF-Radius-Class`, not `cVPN3000-IETF-RADIUS-Class`.

**Table E-2** Security Appliance Supported LDAP Cisco Schema Attributes

| Attribute Name/<br>OID (Object Identifier) | VPN<br>3000 | ASA | PIX | Attr.<br>OID <sup>1</sup> | Syntax/<br>Type | Single<br>or<br>Multi-<br>Valued | Possible Values                                                                                                                                                |
|--------------------------------------------|-------------|-----|-----|---------------------------|-----------------|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cVPN3000-Access-Hours                      | Y           | Y   | Y   | 1                         | String          | Single                           | Name of the time-range (i.e., Business-Hours)                                                                                                                  |
| cVPN3000-Simultaneous-Logins               | Y           | Y   | Y   | 2                         | Integer         | Single                           | 0-2147483647                                                                                                                                                   |
| cVPN3000-Primary-DNS                       | Y           | Y   | Y   | 3                         | String          | Single                           | An IP address                                                                                                                                                  |
| cVPN3000-Secondary-DNS                     | Y           | Y   | Y   | 4                         | String          | Single                           | An IP address                                                                                                                                                  |
| cVPN3000-Primary-WINS                      | Y           | Y   | Y   | 5                         | String          | Single                           | An IP address                                                                                                                                                  |
| cVPN3000-Secondary-WINS                    | Y           | Y   | Y   | 6                         | String          | Single                           | An IP address                                                                                                                                                  |
| cVPN3000-SEP-Card-Assignment               |             |     |     | 7                         | Integer         | Single                           | Not used                                                                                                                                                       |
| cVPN3000-Tunneling-Protocols               | Y           | Y   | Y   | 8                         | Integer         | Single                           | 1 = PPTP<br>2 = L2TP<br>4 = IPsec<br>8 = L2TP/IPsec<br>16 = WebVPN.<br>8 and 4 are mutually exclusive<br>(0 - 11, 16 - 27 are legal values)                    |
| cVPN3000-IPsec-Sec-Association             | Y           |     |     | 9                         | String          | Single                           | Name of the security association                                                                                                                               |
| cVPN3000-IPsec-Authentication              | Y           |     |     | 10                        | Integer         | Single                           | 0 = None<br>1 = RADIUS<br>2 = LDAP (authorization only)<br>3 = NT Domain<br>4 = SDI<br>5 = Internal<br>6 = RADIUS with Expiry<br>7 = Kerberos/Active Directory |
| cVPN3000-IPsec-Banner1                     | Y           | Y   | Y   | 11                        | String          | Single                           | Banner string                                                                                                                                                  |
| cVPN3000-IPsec-Allow-Passwd-Store          | Y           | Y   | Y   | 12                        | Boolean         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                    |
| cVPN3000-Use-Client-Address                | Y           |     |     | 13                        | Boolean         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                    |



**Table E-2** Security Appliance Supported LDAP Cisco Schema Attributes (continued)

| Attribute Name/<br>OID (Object Identifier) | VPN<br>3000 | ASA | PIX | Attr.<br>OID <sup>1</sup> | Syntax/<br>Type | Single<br>or<br>Multi-<br>Valued | Possible Values                                                                                                                           |
|--------------------------------------------|-------------|-----|-----|---------------------------|-----------------|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| cVPN3000-PPTP-Encryption                   | Y           |     |     | 14                        | Integer         | Single                           | Bitmap:<br>1 = Encryption required<br>2 = 40 bits<br>4 = 128 bits<br>8 = Stateless-Required<br>Example: 15 =<br>40/128-Encr/Stateless-Req |
| cVPN3000-L2TP-Encryption                   | Y           |     |     | 15                        | Integer         | Single                           | Bitmap:<br>1 = Encryption required<br>2 = 40 bit<br>4 = 128 bits<br>8 = Stateless-Req<br>15 =<br>40/128-Encr/Stateless-Req                |
| cVPN3000-IPSec-Split-Tunnel-List           | Y           | Y   | Y   | 16                        | String          | Single                           | Specifies the name of the network or access list that describes the split tunnel inclusion list.                                          |
| cVPN3000-IPSec-Default-Domain              | Y           | Y   | Y   | 17                        | String          | Single                           | Specifies the single default domain name to send to the client (1-255 characters).                                                        |
| cVPN3000-IPSec-Split-DNS-Name              | Y           | Y   | Y   | 18                        | String          | Single                           | Specifies the list of secondary domain names to send to the client (1-255 characters).                                                    |
| cVPN3000-IPSec-Tunnel-Type                 | Y           | Y   | Y   | 19                        | Integer         | Single                           | 1 = LAN-to-LAN<br>2 = Remote access                                                                                                       |
| cVPN3000-IPSec-Mode-Config                 | Y           | Y   | Y   | 20                        | Boolean         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                               |
| cVPN3000-IPSec-User-Group-Lock             | Y           |     |     | 21                        | Boolean         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                               |
| cVPN3000-IPSec-Over-UDP                    | Y           | Y   | Y   | 22                        | Boolean         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                               |
| cVPN3000-IPSec-Over-UDP-Port               | Y           | Y   | Y   | 23                        | Integer         | Single                           | 4001 - 49151, default = 10000                                                                                                             |
| cVPN3000-IPSec-Banner2                     | Y           | Y   | Y   | 24                        | String          | Single                           | Banner string                                                                                                                             |
| cVPN3000-PPTP-MPPC-Compression             | Y           |     |     | 25                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                               |

**Table E-2** Security Appliance Supported LDAP Cisco Schema Attributes (continued)

| Attribute Name/<br>OID (Object Identifier)        | VPN<br>3000 | ASA | PIX | Attr.<br>OID <sup>1</sup> | Syntax/<br>Type | Single<br>or<br>Multi-<br>Valued | Possible Values                                                                                                                                                                  |
|---------------------------------------------------|-------------|-----|-----|---------------------------|-----------------|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cVPN3000-L2TP-MPPC-Compression                    | Y           |     |     | 26                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                                      |
| cVPN3000-IPSec-IP-Compression                     | Y           | Y   | Y   | 27                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                                      |
| cVPN3000-IPSec-IKE-Peer-ID-Check                  | Y           | Y   | Y   | 28                        | Integer         | Single                           | 1 = Required<br>2 = If supported by peer<br>certificate<br>3 = Do not check                                                                                                      |
| cVPN3000-IKE-Keep-Alives                          | Y           | Y   | Y   | 29                        | Boolean         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                                      |
| cVPN3000-IPSec-Auth-On-Rekey                      | Y           | Y   | Y   | 30                        | Boolean         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                                      |
| cVPN3000-Required-Client-<br>Firewall-Vendor-Code | Y           | Y   | Y   | 31                        | Integer         | Single                           | 1 = Cisco Systems (with Cisco<br>Integrated Client)<br>2 = Zone Labs<br>3 = NetworkICE<br>4 = Sygate<br>5 = Cisco Systems (with Cisco<br>Intrusion Prevention Security<br>Agent) |

**Table E-2** Security Appliance Supported LDAP Cisco Schema Attributes (continued)

| Attribute Name/<br>OID (Object Identifier)         | VPN<br>3000 | ASA | PIX | Attr.<br>OID <sup>1</sup> | Syntax/<br>Type | Single<br>or<br>Multi-<br>Valued | Possible Values                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------|-------------|-----|-----|---------------------------|-----------------|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cVPN3000-Required-Client-Firewall-Product-Code     | Y           | Y   | Y   | 32                        | Integer         | Single                           | Cisco Systems Products:<br>1 = Cisco Intrusion Prevention Security Agent or Cisco Integrated Client (CIC)<br><br>Zone Labs Products:<br>1 = Zone Alarm<br>2 = Zone AlarmPro<br>3 = Zone Labs Integrity<br><br>NetworkICE Product:<br>1 = BlackIce Defender/Agent<br><br>Sygate Products:<br>1 = Personal Firewall<br>2 = Personal Firewall Pro<br>3 = Security Agent |
| cVPN3000-Required-Client-Firewall-Description      | Y           | Y   | Y   | 33                        | String          | Single                           | String                                                                                                                                                                                                                                                                                                                                                               |
| cVPN3000-Require-Individual-User-Auth              | Y           | Y   | Y   | 34                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                                                                                                                                                                                                                          |
| cVPN3000-Require-HW-Client-Auth                    | Y           | Y   | Y   | 35                        | Boolean         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                                                                                                                                                                                                                          |
| cVPN3000-Authenticated-User-Idle-Timeout           | Y           | Y   | Y   | 36                        | Integer         | Single                           | 1 - 35791394 minutes                                                                                                                                                                                                                                                                                                                                                 |
| cVPN3000-Cisco-IP-Phone-Bypass                     | Y           | Y   | Y   | 37                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                                                                                                                                                                                                                          |
| cVPN3000-IPSec-Split-Tunneling-Policy              | Y           | Y   | Y   | 38                        | Integer         | Single                           | 0 = Tunnel everything<br>1 = Split tunneling<br>2 = Local LAN permitted                                                                                                                                                                                                                                                                                              |
| cVPN3000-IPSec-Required-Client-Firewall-Capability | Y           | Y   | Y   | 39                        | Integer         | Single                           | 0 = None<br>1 = Policy defined by remote FW Are-You-There (AYT)<br>2 = Policy pushed CPP<br>4 = Policy from server                                                                                                                                                                                                                                                   |

**Table E-2** Security Appliance Supported LDAP Cisco Schema Attributes (continued)

| Attribute Name/<br>OID (Object Identifier)     | VPN<br>3000 | ASA | PIX | Attr.<br>OID <sup>1</sup> | Syntax/<br>Type | Single<br>or<br>Multi-<br>Valued | Possible Values                                                                                                                                                                                                                                                                           |
|------------------------------------------------|-------------|-----|-----|---------------------------|-----------------|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cVPN3000-IPSec-Client-Firewall-Filter-Name     | Y           |     |     | 40                        | String          | Single                           | Specifies the name of the filter to be pushed to the client as firewall policy.                                                                                                                                                                                                           |
| cVPN3000-IPSec-Client-Firewall-Filter-Optional | Y           | Y   | Y   | 41                        | Integer         | Single                           | 0 = Required<br>1 = Optional                                                                                                                                                                                                                                                              |
| cVPN3000-IPSec-Backup-Servers                  | Y           | Y   | Y   | 42                        | String          | Single                           | 1 = Use Client-Configured list<br>2 = Disabled and clear client list<br>3 = Use Backup Server list                                                                                                                                                                                        |
| cVPN3000-IPSec-Backup-Server-List              | Y           | Y   | Y   | 43                        | String          | Single                           | Server Addresses (space delimited)                                                                                                                                                                                                                                                        |
| cVPN3000-Client-Intercept-DHCP-Configure-Msg   | Y           | Y   | Y   | 44                        | Boolean         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                                                                                                                                               |
| cVPN3000-MS-Client-Subnet-Mask                 | Y           | Y   | Y   | 45                        | String          | Single                           | An IP address                                                                                                                                                                                                                                                                             |
| cVPN3000-Allow-Network-Extension-Mode          | Y           | Y   | Y   | 46                        | Boolean         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                                                                                                                                               |
| cVPN3000-Strip-Realm                           | Y           | Y   | Y   | 47                        | Boolean         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                                                                                                                                               |
| cVPN3000-Cisco-AV-Pair                         | Y           | Y   | Y   | 48                        | String          | Multi                            | An octet string in the following format:<br><br>[Prefix] [Action] [Protocol]<br>[Source] [Source Wildcard Mask] [Destination]<br>[Destination Wildcard Mask]<br>[Established] [Log] [Operator] [Port]<br><br>For more information, see <a href="#">“Cisco -AV-Pair Attribute Syntax.”</a> |
| cVPN3000-User-Auth-Server-Name                 | Y           |     |     | 49                        | String          | Single                           | IP address or hostname                                                                                                                                                                                                                                                                    |
| cVPN3000-User-Auth-Server-Port                 | Y           |     |     | 50                        | Integer         | Single                           | Port number for server protocol                                                                                                                                                                                                                                                           |
| cVPN3000-User-Auth-Server-Secret               | Y           |     |     | 51                        | String          | Single                           | Server password                                                                                                                                                                                                                                                                           |
| cVPN3000-Confidence-Interval                   | Y           | Y   | Y   | 52                        | Integer         | Single                           | 10 - 300 seconds                                                                                                                                                                                                                                                                          |
| cVPN3000-Cisco-LEAP-Bypass                     | Y           | Y   | Y   | 53                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                                                                                                                                                                                                               |
| cVPN3000-DHCP-Network-Scope                    | Y           | Y   | Y   | 54                        | String          | Single                           | IP address                                                                                                                                                                                                                                                                                |

**Table E-2**      **Security Appliance Supported LDAP Cisco Schema Attributes (continued)**

| Attribute Name/<br>OID (Object Identifier)   | VPN<br>3000 | ASA | PIX | Attr.<br>OID <sup>1</sup> | Syntax/<br>Type | Single<br>or<br>Multi-<br>Valued | Possible Values                                                                                                                                                                                            |
|----------------------------------------------|-------------|-----|-----|---------------------------|-----------------|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cVPN3000-Client-Type-Version-Limiting        | Y           | Y   | Y   | 55                        | String          | Single                           | IPsec VPN client version number string                                                                                                                                                                     |
| cVPN3000-WebVPN-Content-Filter-Parameters    | Y           | Y   |     | 56                        | Integer         | Single                           | 1 = Java & ActiveX<br>2 = Java scripts<br>4 = Images<br>8 = Cookies in images<br>Add the values to filter multiple parameters. For example: enter 10 to filter both Java scripts and cookies. (10 = 2 + 8) |
| cVPN3000-WebVPN-Enable-functions             |             |     |     | 57                        | Integer         | Single                           | Not used - deprecated                                                                                                                                                                                      |
| cVPN3000-WebVPN-Exchange-Server-Address      |             |     |     | 58                        | String          | Single                           | Not used - deprecated                                                                                                                                                                                      |
| cVPN3000-WebVPN-Exchange-Server-NETBIOS-Name |             |     |     | 59                        | String          | Single                           | Not used - deprecated                                                                                                                                                                                      |
| cVPN3000-Port-Forwarding-Name                | Y           | Y   |     | 60                        | String          | Single                           | Name string (for example, "Corporate-Apps")                                                                                                                                                                |
| cVPN3000-IETF-Radius-Framed-IP-Address       | Y           | Y   | Y   | 61                        | String          | Single                           | An IP address                                                                                                                                                                                              |
| cVPN3000-IETF-Radius-Framed-IP-Netmask       | Y           | Y   | Y   | 62                        | String          | Single                           | An IP address                                                                                                                                                                                              |
| cVPN3000-IETF-Radius-Session-Timeout         | Y           | Y   | Y   | 63                        | Integer         | Single                           | 1-35791394 minutes<br>0 = Unlimited                                                                                                                                                                        |
| cVPN3000-IETF-Radius-Idle-Timeout            | Y           | Y   | Y   | 64                        | Integer         | Single                           | 1-35791394 minutes<br>0 = Unlimited                                                                                                                                                                        |
| cVPN3000-IETF-Radius-Class                   | Y           | Y   | Y   | 65                        | String          | Single                           | Group name string. Use any of the these three formats:<br>OU=Engineering<br>OU=Engineering;<br>Engineering                                                                                                 |
| cVPN3000-IETF-Radius-Filter-Id               | Y           | Y   | Y   | 66                        | String          | Single                           | An access-list                                                                                                                                                                                             |
| cVPN3000-Authorization-Required              | Y           |     |     | 67                        | Integer         | Single                           | 0 = No<br>1 = Yes                                                                                                                                                                                          |
| cVPN3000-Authorization-Type                  | Y           |     |     | 68                        | Integer         | Single                           | 0 = None<br>1 = RADIUS<br>2 = LDAP                                                                                                                                                                         |

**Table E-2 Security Appliance Supported LDAP Cisco Schema Attributes (continued)**

| Attribute Name/<br>OID (Object Identifier)                | VPN<br>3000 | ASA | PIX | Attr.<br>OID <sup>1</sup> | Syntax/<br>Type | Single<br>or<br>Multi-<br>Valued | Possible Values                                                                                  |
|-----------------------------------------------------------|-------------|-----|-----|---------------------------|-----------------|----------------------------------|--------------------------------------------------------------------------------------------------|
| cVPN3000-DN-Field                                         | Y           | Y   | Y   | 69                        | String          | Single                           | Possible values: UID, OU, O, CN, L, SP, C, EA, T, N, GN, SN, I, GENQ, DNQ, SER, use-entire-name. |
| cVPN3000-WebVPN-URL-List                                  |             | Y   |     | 70                        | String          | Single                           | URL-list name                                                                                    |
| cVPN3000-WebVPN-Forwarded-Ports                           |             | Y   |     | 71                        | String          | Single                           | Port-Forward list name                                                                           |
| cVPN3000-WebVPN-ACL-Filters                               |             | Y   |     | 72                        | String          | Single                           | Access-List name                                                                                 |
| cVPN3000-WebVPN-Homepage                                  | Y           | Y   |     | 73                        | String          | Single                           | A url such as<br>http://example-portal.com.                                                      |
| cVPN3000-WebVPN-Single-Sign-On-<br>Server-Name            |             | Y   |     | 74                        | String          | Single                           | Name of the SSO Server (1 - 31<br>chars)                                                         |
| cVPN3000-WebVPN-URL-Entry-Enable                          | Y           | Y   |     | 75                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |
| cVPN3000-WebVPN-File-Access-Enable                        | Y           | Y   |     | 76                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |
| cVPN3000-WebVPN-File-Server-Entry-<br>Enable              | Y           | Y   |     | 77                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |
| cVPN3000-WebVPN-File-Server-Browsing-<br>Enable           | Y           | Y   |     | 78                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |
| cVPN3000-WebVPN-Port-Forwarding-<br>Enable                | Y           | Y   |     | 79                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |
| cVPN3000-WebVPN-Port-Forwarding-<br>Exchange-Proxy-Enable | Y           | Y   |     | 80                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |
| cVPN3000-WebVPN-Port-Forwarding-<br>HTTP-Proxy-Enable     | Y           | Y   |     | 81                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |
| cVPN3000-WebVPN-Port-Forwarding-<br>Auto-Download-Enable  | Y           | Y   |     | 82                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |
| cVPN3000-WebVPN-Citrix-Support-Enable                     | Y           | Y   |     | 83                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |
| cVPN3000-WebVPN-Apply-ACL-Enable                          | Y           | Y   |     | 84                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |
| cVPN3000-WebVPN-SVC-Enable                                | Y           | Y   |     | 85                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |
| cVPN3000-WebVPN-SVC-Required-Enable                       | Y           | Y   |     | 86                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                                                      |

**Table E-2** Security Appliance Supported LDAP Cisco Schema Attributes (continued)

| Attribute Name/<br>OID (Object Identifier) | VPN<br>3000 | ASA | PIX | Attr.<br>OID <sup>1</sup> | Syntax/<br>Type | Single<br>or<br>Multi-<br>Valued | Possible Values                                                         |
|--------------------------------------------|-------------|-----|-----|---------------------------|-----------------|----------------------------------|-------------------------------------------------------------------------|
| cVPN3000-WebVPN-SVC-Keep-Enable            | Y           | Y   |     | 87                        | Integer         | Single                           | 0 = Disabled<br>1 = Enabled                                             |
| cVPN3000-IE-Proxy-Server                   | Y           |     |     | 88                        | String          | Single                           | IP address                                                              |
| cVPN3000-IE-Proxy-Method                   | Y           |     |     | 89                        | Integer         | Single                           | 1 = No Modify<br>2 = No Proxy<br>3 = Auto Detect<br>4 = Other           |
| cVPN3000-IE-Proxy-Exception-List           | Y           |     |     | 90                        | String          | Single                           | newline (\n)-separated list of<br>DNS domains                           |
| cVPN3000-IE-Proxy-Bypass-Local             | Y           |     |     | 91                        | Integer         | Single                           | 0 = None<br>1 = Local                                                   |
| cVPN3000-Tunnel-Group-Lock                 |             | Y   | Y   | 92                        | String          | Single                           | Name of the tunnel group or<br>"none"                                   |
| cVPN3000-Firewall-ACL-In                   |             | Y   | Y   | 93                        | String          | Single                           | Access list ID                                                          |
| cVPN3000-Firewall-ACL-Out                  |             | Y   | Y   | 94                        | String          | Single                           | Access list ID                                                          |
| cVPN3000-PFS-Required                      | Y           | Y   | Y   | 95                        | Boolean         | Single                           | 0 = No<br>1 = Yes                                                       |
| cVPN3000-WebVPN-SVC-Keepalive              | Y           | Y   |     | 96                        | Integer         | Single                           | 0 = Disabled<br>n = Keepalive value in seconds<br>(15 - 600)            |
| cVPN3000-WebVPN-SVC-Client-DPD             | Y           | Y   |     | 97                        | Integer         | Single                           | 0 = Disabled<br>n = Dead Peer Detection value<br>in seconds (30 - 3600) |
| cVPN3000-WebVPN-SVC-Gateway-DPD            | Y           | Y   |     | 98                        | Integer         | Single                           | 0 = Disabled<br>n = Dead Peer Detection value<br>in seconds (30 - 3600) |
| cVPN3000-WebVPN-SVC-Rekey-Period           | Y           | Y   |     | 99                        | Integer         | Single                           | 0 = Disabled<br>n = Retry period in minutes (4 -<br>10080)              |
| cVPN3000-WebVPN-SVC-Rekey-Method           | Y           | Y   |     | 100                       | Integer         | Single                           | 0 = None<br>1 = SSL<br>2 = New tunnel<br>3 = Any (sets to SSL)          |
| cVPN3000-WebVPN-SVC-Compression            | Y           | Y   |     | 101                       | Integer         | Single                           | 0 = None<br>1 = Deflate Compression                                     |

- To get the complete Object Identifier of each attribute, append the number in the column to the end of 1.2.840.113556.8000.795.2. Thus, the OID of the first attribute in the table, cVPN3000-Access-Hours, is 1.2.840.113556.8000.795.2.1. Likewise, the OID of the last attribute in the table, cVPN3000-WebVPN-SVC-Compression, is 1.2.840.113556.8000.795.2.115.

## Cisco -AV-Pair Attribute Syntax

The syntax of each Cisco-AV-Pair rule is as follows:

[Prefix] [Action] [Protocol] [Source] [Source Wildcard Mask] [Destination] [Destination Wildcard Mask] [Established] [Log] [Operator] [Port]:

| Field                     | Description                                                                                                                                                                                                                    |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Prefix                    | A unique identifier for the AV pair. For example: <code>ip:inacl#1=</code> (used for standard ACLs) or <code>webvpn:inacl#</code> (used for WebVPN ACLs). This field only appears when the filter has been sent as an AV pair. |
| Action                    | Action to perform if rule matches: deny, permit.                                                                                                                                                                               |
| Protocol                  | Number or name of an IP protocol. Either an integer in the range 0-255 or one of the following keywords: icmp, igmp, ip, tcp, udp.                                                                                             |
| Source                    | Network or host that sends the packet. It is specified as an IP address, a hostname, or the keyword “any”. If specified as an IP address, the source wildcard mask must follow.                                                |
| Source Wildcard Mask      | The wildcard mask applied to the source address.                                                                                                                                                                               |
| Destination               | Network or host that receives the packet. It is specified as an IP address, a hostname, or the keyword “any”. If specified as an IP address, the source wildcard mask must follow.                                             |
| Destination Wildcard Mask | The wildcard mask applied to the destination address.                                                                                                                                                                          |
| Log                       | Generates a FILTER log message. You must use this keyword to generate events of severity level 9.                                                                                                                              |
| Operator                  | Logic operators: greater than, less than, equal to, not equal to.                                                                                                                                                              |
| Port                      | The number of a TCP or UDP port in the range 0-65535.                                                                                                                                                                          |

For example:

```
ip:inacl#1=deny ip 10.155.10.0 0.0.0.255 10.159.2.0 0.0.0.255 log
ip:inacl#2=permit TCP any host 10.160.0.1 eq 80 log

webvpn:inacl#1=permit url http://www.cnn.com
webvpn:inacl#2=deny smtp any host 10.1.3.5
webvpn:inacl#3=permit url cifs://mar_server/peopleshare1
```



**Note**

- Use Cisco-AV pair entries with the ip:inacl# prefix to enforce ACLs for remote IPsec and SSL VPN Client (SVC) tunnels.
- Use Cisco-AV pair entries with the webvpn:inacl# prefix to enforce ACLs for WebVPN clientless (browser-mode) tunnels.

Table E-3 lists the tokens for the Cisco-AV-Pair attribute:

**Table E-3 Security Appliance-Supported Tokens**

| Token             | Syntax Field     | Description                                                                                                                                         |
|-------------------|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| ip:inacl#Num=     | N/A (Identifier) | (Where <i>Num</i> is a unique integer.) Starts all AV pair access control lists. Enforces ACLs for remote IPsec and SSL VPN (SVC) tunnels.          |
| webvpn:inacl#Num= | N/A (Identifier) | (Where <i>Num</i> is a unique integer.) Starts all WebVPN AV pair access control lists. Enforces ACLs for WebVPN clientless (browser-mode) tunnels. |
| deny              | Action           | Denies action. (Default)                                                                                                                            |
| permit            | Action           | Allows action.                                                                                                                                      |
| icmp              | Protocol         | Internet Control Message Protocol (ICMP)                                                                                                            |
| 1                 | Protocol         | Internet Control Message Protocol (ICMP)                                                                                                            |
| IP                | Protocol         | Internet Protocol (IP)                                                                                                                              |
| 0                 | Protocol         | Internet Protocol (IP)                                                                                                                              |
| TCP               | Protocol         | Transmission Control Protocol (TCP)                                                                                                                 |
| 6                 | Protocol         | Transmission Control Protocol (TCP)                                                                                                                 |
| UDP               | Protocol         | User Datagram Protocol (UDP)                                                                                                                        |
| 17                | Protocol         | User Datagram Protocol (UDP)                                                                                                                        |
| any               | Hostname         | Rule applies to any host.                                                                                                                           |
| host              | Hostname         | Any alpha-numeric string that denotes a hostname.                                                                                                   |
| log               | Log              | When the event is hit, a filter log message appears. (Same as permit and log or deny and log.)                                                      |
| lt                | Operator         | Less than value                                                                                                                                     |
| gt                | Operator         | Greater than value                                                                                                                                  |
| eq                | Operator         | Equal to value                                                                                                                                      |
| neq               | Operator         | Not equal to value                                                                                                                                  |
| range             | Operator         | Inclusive range. Should be followed by two values.                                                                                                  |

## Example Security Appliance Authorization Schema

This section provides a sample of an LDAP schema. This schema supports the security appliance class and attributes. It is specific to the Microsoft Active Directory LDAP server. Use it as a model, with [Table E-2](#), to define your own schema for your own LDAP server.

## Schema 3k\_schema.ldif

```

dn:
CN=cVPN3000-Access-Hours,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation,DC=com
changetype: add
adminDisplayName: cVPN3000-Access-Hours
attributeID: 1.2.840.113556.1.8000.795.2.1
attributeSyntax: 2.5.5.3
cn: cVPN3000-Access-Hours
instanceType: 4
isSingleValued: TRUE
LDAPDisplayName: cVPN3000-Access-Hours
distinguishedName:

CN=cVPN3000-Access-Hours,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation,DC=com
objectCategory:
 CN=Attribute-Schema,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation,DC=com
objectClass: attributeSchema
oMSyntax: 27
name: cVPN3000-Access-Hours
showInAdvancedViewOnly: TRUE

....
.... (define subsequent security appliance authorization attributes here)
....

dn:
CN=cVPN3000-Primary-DNS,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation,DC=com
changetype: add
adminDisplayName: cVPN3000-Primary-DNS
attributeID: 1.2.840.113556.1.8000.795.2.3
attributeSyntax: 2.5.5.3
cn: cVPN3000-Primary-DNS
instanceType: 4
isSingleValued: TRUE
LDAPDisplayName: cVPN3000-Primary-DNS
distinguishedName:
CN=cVPN3000-Primary-DNS,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation,DC=com
objectCategory:
 CN=Attribute-Schema,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation,DC=com
objectClass: attributeSchema
oMSyntax: 27
name: cVPN3000-Primary-DNS
showInAdvancedViewOnly: TRUE

....
.... (define subsequent security appliance authorization attributes here)
....

dn:
CN=cVPN3000-Confidence-Interval,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation
,DC=com
changetype: add
adminDisplayName: cVPN3000-Confidence-Interval
attributeID: 1.2.840.113556.1.8000.795.2.52
attributeSyntax: 2.5.5.9
cn: cVPN3000-Confidence-Interval
instanceType: 4
isSingleValued: TRUE
LDAPDisplayName: cVPN3000-Confidence-Interval
distinguishedName:

```

```

CN=cVPN3000-Confidence-Interval,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation
,DC=com
objectCategory:

DN:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-

dn:
CN=cVPN3000-User-Authorization,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation,
DC=com
changetype: add
adminDisplayName: cVPN3000-User-Authorization
adminDescription: Cisco Class Schema
cn: cVPN3000-User-Authorization
defaultObjectCategory:

CN=cVPN3000-User-Authorization,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation,
DC=com
defaultSecurityDescriptor:
 D: (A;;RPWPCRCDCCLCLOLORCWOWSDDDTDTSW;;;DA) (A;;RPWPCRCDCCLCLOLORCWOWSDDDTDTSW;;;SY)
 (A;;RPLCLORC;;;AU)
governsID: 1.2.840.113556.1.8000.795.1.1
instanceType: 4
LDAPDisplayName: cVPN3000-User-Authorization

mustContain: cn
mayContain: cVPN3000-Access-Hours
mayContain: cVPN3000-Simultaneous-Logins
mayContain: cVPN3000-Primary-DNS
...
mayContain: cVPN3000-Confidence-Interval
mayContain: cVPN3000-Cisco-LEAP-Bypass

distinguishedName:

CN=cVPN3000-User-Authorization,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation,
DC=com
objectCategory:
 CN=Class-Schema,CN=Schema,CN=Configuration,OU=People,DC=ExampleCorporation,DC=com
objectClass: classSchema
objectClassCategory: 1
possSuperiors: organizationalUnit
name: cVPN3000-User-Authorization
rDNAttID: cn
showInAdvancedViewOnly: TRUE
subClassOf: top
systemOnly: FALSE

DN:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-
systemOnly: FALSE

DN:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-

```

## Loading the Schema in the LDAP Server

**Note**

The directions in this section are specific to the Microsoft Active Directory LDAP server. If you have a different type of server, refer to your server documentation for information on loading a schema.

To load the schema on the LDAP server, enter the following command from the directory where the schema file resides: `ldifde -i -f Schema Name`. For example: `ldifde -i -f 3k_schema.ldif`

## Defining User Permissions

**Note**

The directions in this section are specific to the Microsoft Active Directory LDAP server. If you have a different type of server, see your server documentation to define and load user attributes.

For each user authorizing to your LDAP server, define a user file. A user file defines all the security appliance attributes and values associated with a particular user. Each user is an object of the class `cVPN3000-User-Authorization`. To define the user file, use any text editor. The file must have the extension `.ldif`. (For an example user file, see [Robin.ldif](#).)

To load the user file on the LDAP server, enter the following command on the directory where your version of the `ldap_user.ldif` file resides: `ldifde -i -f ldap_user.ldif`. For example: `ldifde -i -f Robin.ldif`

After you have created and loaded both the schema and the user file, your LDAP server is ready to process security appliance authorization requests.

## Example User File

This section provides a sample user file for the user Robin.

### Robin.ldif

```
dn: cn=Robin,OU=People,DC=ExampleCorporation,DC=com
changetype: add
cn: Robin
CVPN3000-Access-Hours: Corporate_time
cVPN3000-Simultaneous-Logins: 2
cVPN3000-IPSec-Over-UDP: TRUE
CVPN3000-IPSec-Over-UDP-Port: 12125
cVPN3000-IPSec-Banner1: Welcome to the Example Corporation!!!
cVPN3000-IPSec-Banner2: Unauthorized access is prohibited!!!!
cVPN3000-Primary-DNS: 10.10.4.5
CVPN3000-Secondary-DNS: 10.11.12.7
CVPN3000-Primary-WINS: 10.20.1.44
CVPN3000-SEP-Card-Assignment: 1
CVPN3000-IPSec-Tunnel-Type: 2
CVPN3000-Tunneling-Protocols: 7
cVPN3000-Confidence-Interval: 300
cVPN3000-IPSec-Allow-Passwd-Store: TRUE
objectClass: cVPN3000-User-Authorization
```

## Reviewing Examples of Active Directory Configurations

This section presents example procedures for configuring authentication and authorization on the security appliance using the Microsoft Active Directory server. It includes the following topics:

- [Example 1: Configuring LDAP Authorization with Microsoft Active Directory \(ASA/PIX\)](#)
- [Example 2: Configuring LDAP Authentication with Microsoft Active Directory](#)
- [Example 3: LDAP Authentication and LDAP Authorization with Microsoft Active Directory](#)

### Example 1: Configuring LDAP Authorization with Microsoft Active Directory (ASA/PIX)

This example presents a configuration procedure for authentication using SDI and authorization using LDAP and Microsoft Active Directory (AD). To execute this sample procedure, perform the following steps:

- Step 1** Using LDIF files, create the cVPN3000-User-Authorization record on the Microsoft AD database. This record contains the Cisco VPN authorization attributes for the user.




**Note** Contact Cisco TAC to obtain the Microsoft AD LDAP schema for Cisco VPN attributes.

- Step 2** To confirm the new record, choose **Start > Settings > Control Panel > Administrative Tools > Active Directory Users and Computers**.

The Active Directory Users and Computers window appears as shown in [Figure E-2](#).

**Figure E-2** Active Directory Users and Computers Window



- Step 3** On the security appliance, create a aaa-server record for the LDAP server. For this example, these authorization records are stored in the Franklin-Altiga folder. The necessary steps are shown in the following commands:

```
hostname(config)# aaa-server ldap-authorize-grp protocol ldap
hostname(config-aaa-server-group)# aaa-server ldap-authorize-grp host 10.1.1.4
hostname(config-aaa-server-host)# ldap-base-dn ou=Franklin-Altiga,dc=frdevtestad,dc=local
hostname(config-aaa-server-host)# ldap-scope subtree
```

```

hostname(config-aaa-server-host) # ldap-naming-attribute cn
hostname(config-aaa-server-host) # ldap-login-password anypassword
hostname(config-aaa-server-host) # ldap-login-dn cn=Administrator,cn=Users,
dc=frdevtestad,dc=local
hostname(config-aaa-server-host) # ldap-attribute-map LdapSvrName
hostname(config-aaa-server-host) #

```

**Step 4** Create a tunnel group that specifies SDI Authentication and LDAP authorization, as shown in the following example commands:

```

hostname(config) # tunnel-group ipsec-tunnelgroup type ipsec-ra
hostname(config) # tunnel-group ipsec-tunnelgroup general-attributes
hostname(config) # authentication-server-group sdi-group
hostname(config) # authorization-server-group ldap-authorize-group
hostname(config) #

```



**Note**


This example does not show the configuration for sdi-group.

## Example 2: Configuring LDAP Authentication with Microsoft Active Directory

This example presents a configuration procedure for LDAP authentication with Microsoft Active Directory. To secure the user credentials during transmission, this procedure configures the security appliance to exchange messages with the LDAP directory over a SSL connection. It also configures the security appliance to interpret the department attribute in the Microsoft AD user record as the group policy to which the user is assigned. The authorization attributes for this group are retrieved from a RADIUS server.


View the user records by clicking the User folder in the Active Directory Users and Computers window as shown in [Figure E-3](#).

**Figure E-3** Active Directory Users and Computers Window Showing User Folder



Review specific user attributes and values by right-clicking the username and clicking Properties. The Username Properties dialog box appears as shown in Figure E-4.

**Figure E-4** The Username Properties Dialog Box



**Note**

The department attribute is configured under the Organization tab in the Active Directory Users and Computers window.

To configure this example, perform the following steps on the security appliance:

- Step 1** Create a aaa-server record for the LDAP authentication server and use the ldap-base-dn to specify the search location for the Active Directory user records as shown in the following example commands:

```
hostname(config)# aaa-server ldap-authenticate-grp protocol ldap
hostname(config-aaa-server-group)# aaa-server ldap-authenticate-grp host 10.1.1.4
hostname(config-aaa-server-host)# ldap-base-dn cn=Users,dc=frdevtestad,dc=local
hostname(config-aaa-server-host)# ldap-scope subtree
hostname(config-aaa-server-host)# ldap-naming-attribute cn
hostname(config-aaa-server-host)# ldap-login-password anypassword
hostname(config-aaa-server-host)# ldap-login-dn cn=Administrator,cn=Users,dc=frdevtestad,dc=local
hostname(config-aaa-server-host)#
```

- Step 2** Create an LDAP mapping table entry to map the AD attribute department to the Cisco attribute cVPN3000-IETF-Radius-Class as shown in the following example commands:

```
hostname(config)# ldap attribute-map ActiveDirectoryMapTable
```

```
hostname(config-ldap-attribute-map)# map-name department cVPN3000-IETF-Radius-Class
hostname(config-ldap-attribute-map)#
```

**Step 3** Configure the name of the LDAP attribute map as shown in the following example command:

```
hostname(config-aaa-server-host)# ldap-attribute-map ActiveDirectoryMapTable
hostname(config-aaa-server-host)#
```

**Step 4** Specify a secure LDAP connection as follows:

```
hostname(config-aaa-server-host)# ldap-over-ssl enable
hostname(config-aaa-server-host)#
```

**Step 5** Create an external group policy that associates the group-name with the RADIUS server. In this example, the user is assigned to the group Engineering as shown in the following example command:

```
hostname(config-aaa-server-host)# group-policy Engineering external server-group
radius-group password anypassword
hostname(config-aaa-server-host)#
```

**Step 6** Create a tunnel group that specifies LDAP authentication as shown in the following example commands:

```
hostname(config)# tunnel-group ipsec-tunnelgroup type ipsec-ra
hostname(config)# tunnel-group ipsec-tunnelgroup general-attributes
hostname(config-tunnel-general)# authentication-server-group ldap-authenticate-grp
hostname(config-tunnel-general)#
```



**Note**

---

The configuration for radius-group is not shown in this example.


---

### Example 3: LDAP Authentication and LDAP Authorization with Microsoft Active Directory

This example presents the procedure for configuring both authentication and authorization using LDAP and Microsoft Active Directory. In the Microsoft user record, the department attribute is interpreted as the group-name for the user. The authorization attributes for this group-name are retrieved from the Active Directory server.

The department attribute is configured under the Organization tab in the Active Directory Users and Computers dialog box as shown in [Figure E-5](#).



**Figure E-5** The Organization Tab of the Active Directory Users and Computer Dialog

To configure this example, perform the following steps on the security appliance:

- Step 1** Create an LDAP mapping table entry to map the Active Directory attribute department to the Cisco attribute cVPN3000-IETF-Radius-Class as shown in the following example commands:
- ```
hostname(config)# ldap attribute-map ActiveDirectoryMapTable
hostname(config-ldap-attribute-map)# map-name department cVPN3000-IETF-Radius-Class
```
- Step 2** Create a aaa-server record for the LDAP authentication server and use the ldap-base-dn to specify the search location for the Active Directory user records as shown in the following example commands:
- ```
hostname(config)# aaa-server ldap-authenticate protocol ldap
hostname(config-aaa-server-group)# aaa-server ldap-authenticate host 10.1.1.4
hostname(config-aaa-server-host)# ldap-base-dn cn=Users,dc=frdevtestad,dc=local
hostname(config-aaa-server-host)# ldap-scope subtree
hostname(config-aaa-server-host)# ldap-naming-attribute cn
hostname(config-aaa-server-host)# ldap-login-password anypassword
hostname(config-aaa-server-host)# ldap-login-dn cn=Administrator,cn=Users,dc=frdevtestad,dc=local
hostname(config-aaa-server-host)#
```
- Step 3** Configure the name of the LDAP attribute map as shown in the following example command:
- ```
hostname(config-aaa-server-host)# ldap-attribute-map ActiveDirectoryMapTable
hostname(config-aaa-server-host)#
```

Step 4 Specify a secure LDAP connection as follows:

```
hostname(config-aaa-server-host)# ldap-over-ssl enable  
hostname(config-aaa-server-host)#
```

Step 5 Create an aaa-server record to configure the LDAP authorization server and use the ldap-base-dn to specify the search location for the Cisco cVPN3000-User-Authorization records as shown in the following example commands:

```
hostname(config-aaa-server-host)# aaa-server ldap-authorize protocol ldap  
hostname(config-aaa-server-host)# aaa-server ldap-authorize host 10.1.1.4  
hostname(config-aaa-server-host)# ldap-base-dn ou=Franklin-Altiga,dc=frdevtestad, dc=local  
hostname(config-aaa-server-host)# ldap-scope subtree  
hostname(config-aaa-server-host)# ldap-naming-attribute cn  
hostname(config-aaa-server-host)# ldap-login-password anypassword  
hostname(config-aaa-server-host)# ldap-login-dn cn=Administrator,cn=Users,  
dc=frdevtestad,dc=local  
hostname(config-aaa-server-host)#
```

Step 6 Create an external group policy that associates the group-name with the LDAP authorization server. In this example, the user is assigned to the group Engineering as shown in the following command:

```
hostname(config-aaa-server-host)# group-policy engineering external server-group  
ldap-authorize  
hostname(config-aaa-server-host)#
```

Step 7 Create a tunnel group that specifies LDAP authentication as shown in the following example commands:

```
hostname(config)# tunnel-group ipsec-tunnelgroup type ipsec-ra  
hostname(config)# tunnel-group ipsec-tunnelgroup general-attributes  
hostname(config-tunnel-general)# authentication-server-group ldap-authenticate  
hostname(config-tunnel-general)#
```

Configuring an External RADIUS Server

This section presents an overview of the RADIUS configuration procedure and defines the Cisco RADIUS and TACACS+ attributes. It includes the following topics:

- [Reviewing the RADIUS Configuration Procedure](#)
- [Security Appliance RADIUS Authorization Attributes](#)
- [Security Appliance TACACS+ Attributes](#)

Reviewing the RADIUS Configuration Procedure

This section describes the RADIUS configuration steps required to support authentication and authorization of the security appliance users. Follow the steps below to set up the RADIUS server to interoperate with the security appliance.

Step 1 Load the security appliance attributes into the RADIUS server. The method you use to load the attributes depends on which type of RADIUS server you are using:

- If you are using Cisco ACS: the server already has these attributes integrated. You can skip this step.

- If you are using a FUNK RADIUS server: Cisco supplies a dictionary file that contains all the security appliance attributes. Obtain this dictionary file, `cisco3k.dct`, from Software Center on CCO or from the security appliance CD-ROM. Load the dictionary file on your server.
- For other vendors' RADIUS servers (for example, Microsoft Internet Authentication Service): you must manually define each security appliance attribute. To define an attribute, use the attribute name or number, type, value, and vendor code (3076). For a list of security appliance RADIUS authorization attributes and values, see [Table E-4](#).

Step 2 Set up the users or groups with the permissions and attributes to send during IPSec/WebVPN tunnel establishment. The permissions or attributes might include access hours, primary DNS, banner, and so forth.

Security Appliance RADIUS Authorization Attributes



Note

Authorization refers to the process of enforcing permissions or attributes. A RADIUS server defined as an authentication server enforces permissions or attributes if they are configured.

[Table E-4](#) lists all the possible security appliance supported attributes that can be used for user authorization.

Table E-4 Security Appliance Supported RADIUS Attributes and Values

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/Type	Single or Multi-Valued	Description or Value
Access-Hours	Y	Y	Y	1	String	Single	Name of the time range, for example, Business-hours
Simultaneous-Logins	Y	Y	Y	2	Integer	Single	An integer 0 to 2147483647
Primary-DNS	Y	Y	Y	5	String	Single	An IP address
Secondary-DNS	Y	Y	Y	6	String	Single	An IP address
Primary-WINS	Y	Y	Y	7	String	Single	An IP address
Secondary-WINS	Y	Y	Y	8	String	Single	An IP address
SEP-Card-Assignment				9	Integer	Single	Not used
Tunneling-Protocols	Y	Y	Y	11	Integer	Single	1 = PPTP 2 = L2TP 4 = IPSec 8 = L2TP/IPSec 16 = WebVPN 4 and 8 are mutually exclusive; 0-11 and 16-27 are legal values.
IPSec-Sec-Association	Y			12	String	Single	Name of the security association

Table E-4 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/Type	Single or Multi-Valued	Description or Value
IPSec-Authentication	Y			13	Integer	Single	0 = None 1 = RADIUS 2 = LDAP (authorization only) 3 = NT Domain 4 = SDI 5 = Internal 6 = RADIUS with Expiry 7 = Kerberos/Active Directory
Banner1	Y	Y	Y	15	String	Single	Banner string
IPSec-Allow-Passwd-Store	Y	Y	Y	16	Boolean	Single	0 = Disabled 1 = Enabled
Use-Client-Address	Y			17	Boolean	Single	0 = Disabled 1 = Enabled
PPTP-Encryption	Y			20	Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bits 4 = 128 bits 8 = Stateless-Required 15 = 40/128-Encr/Stateless-Req
L2TP-Encryption	Y			21	Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bit 4 = 128 bits 8 = Stateless-Req 15 = 40/128-Encr/Stateless-Req
IPSec-Split-Tunnel-List	Y	Y	Y	27	String	Single	Specifies the name of the network/access list that describes the split tunnel inclusion list
IPSec-Default-Domain	Y	Y	Y	28	String	Single	Specifies the single default domain name to send to the client (1-255 characters)

Table E-4 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/Type	Single or Multi-Valued	Description or Value
IPSec-Split-DNS-Names	Y	Y	Y	29	String	Single	Specifies the list of secondary domain names to send to the client (1-255 characters)
IPSec-Tunnel-Type	Y	Y	Y	30	Integer	Single	1 = LAN-to-LAN 2 = Remote access
IPSec-Mode-Config	Y	Y	Y	31	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-User-Group-Lock	Y			33	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-Over-UDP	Y	Y	Y	34	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-Over-UDP-Port	Y	Y	Y	35	Integer	Single	4001 - 49151, default = 10000
Banner2	Y	Y	Y	36	String	Single	A banner string. Banner2 string is concatenated to Banner1 string if configured.
PPTP-MPPC-Compression	Y			37	Integer	Single	0 = Disabled 1 = Enabled
L2TP-MPPC-Compression	Y			38	Integer	Single	0 = Disabled 1 = Enabled
IPSec-IP-Compression	Y	Y	Y	39	Integer	Single	0 = Disabled 1 = Enabled
IPSec-IKE-Peer-ID-Check	Y	Y	Y	40	Integer	Single	1 = Required 2 = If supported by peer certificate 3 = Do not check
IKE-Keep-Alives	Y	Y	Y	41	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-Auth-On-Rekey	Y	Y	Y	42	Boolean	Single	0 = Disabled 1 = Enabled

Table E-4 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/ Type	Single or Multi-Valued	Description or Value
Required-Client- Firewall-Vendor-Code	Y	Y	Y	45	Integer	Single	1 = Cisco Systems (with Cisco Integrated Client) 2 = Zone Labs 3 = NetworkICE 4 = Sygate 5 = Cisco Systems (with Cisco Intrusion Prevention Security Agent)
Required-Client-Firewall-Product-Code	Y	Y	Y	46	Integer	Single	Cisco Systems Products: 1 = Cisco Intrusion Prevention Security Agent or Cisco Integrated Client (CIC) Zone Labs Products: 1 = Zone Alarm 2 = Zone AlarmPro 3 = Zone Labs Integrity NetworkICE Product: 1 = BlackIce Defender/Agent Sygate Products: 1 = Personal Firewall 2 = Personal Firewall Pro 3 = Security Agent
Required-Client-Firewall-Description	Y	Y	Y	47	String	Single	String
Require-HW-Client-Auth	Y	Y	Y	48	Boolean	Single	0 = Disabled 1 = Enabled
Required-Individual-User-Auth	Y	Y	Y	49	Integer	Single	0 = Disabled 1 = Enabled
Authenticated-User-Idle-Timeout	Y	Y	Y	50	Integer	Single	1-35791394 minutes
Cisco-IP-Phone-Bypass	Y	Y	Y	51	Integer	Single	0 = Disabled 1 = Enabled
IPSec-Split-Tunneling-Policy	Y	Y	Y	55	Integer	Single	0 = No split tunneling 1 = Split tunneling 2 = Local LAN permitted

Table E-4 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/Type	Single or Multi-Valued	Description or Value
IPSec-Required-Client-Firewall-Capability	Y	Y	Y	56	Integer	Single	0 = None 1 = Policy defined by remote FW Are-You-There (AYT) 2 = Policy pushed CPP 4 = Policy from server
IPSec-Client-Firewall-Filter-Name	Y			57	String	Single	Specifies the name of the filter to be pushed to the client as firewall policy
IPSec-Client-Firewall-Filter-Optional	Y	Y	Y	58	Integer	Single	0 = Required 1 = Optional
IPSec-Backup-Servers	Y	Y	Y	59	String	Single	1 = Use Client-Configured list 2 = Disable and clear client list 3 = Use Backup Server list
IPSec-Backup-Server-List	Y	Y	Y	60	String	Single	Server Addresses (space delimited)
DHCP-Network-Scope	Y	Y	Y	61	String	Single	IP Address
Intercept-DHCP-Configure-Msg	Y	Y	Y	62	Boolean	Single	0 = Disabled 1 = Enabled
MS-Client-Subnet-Mask	Y	Y	Y	63	Boolean	Single	An IP address
Allow-Network-Extension-Mode	Y	Y	Y	64	Boolean	Single	0 = Disabled 1 = Enabled
Authorization-Type	Y	Y	Y	65	Integer	Single	0 = None 1 = RADIUS 2 = LDAP
Authorization-Required	Y			66	Integer	Single	0 = No 1 = Yes
Authorization-DN-Field	Y	Y	Y	67	String	Single	Possible values: UID, OU, O, CN, L, SP, C, EA, T, N, GN, SN, I, GENQ, DNQ, SER, use-entire-name
IKE-KeepAlive-Confidence-Interval	Y	Y	Y	68	Integer	Single	10-300 seconds
WebVPN-Content-Filter-Parameters	Y	Y		69	Integer	Single	1 = Java ActiveX 2 = Java Script 4 = Image 8 = Cookies in images

Table E-4 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/Type	Single or Multi-Valued	Description or Value
WebVPN-URL-List		Y		71	String	Single	URL-List name
WebVPN-Port-Forward-List		Y		72	String	Single	Port-Forward list name
WebVPN-Access-List		Y		73	String	Single	Access-List name
Cisco-LEAP-Bypass	Y	Y	Y	75	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Homepage	Y	Y		76	String	Single	A URL such as http://example-portal.com
Client-Type-Version-Limiting	Y	Y	Y	77	String	Single	IPsec VPN version number string
WebVPN-Port-Forwarding-Name	Y	Y		79	String	Single	String name (example, "Corporate-Apps"). This text replaces the default string, "Application Access," on the WebVPN home page.
IE-Proxy-Server	Y			80	String	Single	IP address
IE-Proxy-Server-Policy	Y			81	Integer	Single	1 = No Modify 2 = No Proxy 3 = Auto detect 4 = Use Concentrator Setting
IE-Proxy-Exception-List	Y			82	String	Single	newline (\n) separated list of DNS domains
IE-Proxy-Bypass-Local	Y			83	Integer	Single	0 = None 1 = Local
IKE-Keepalive-Retry-Interval	Y	Y	Y	84	Integer	Single	2 - 10 seconds
Tunnel-Group-Lock		Y	Y	85	String	Single	Name of the tunnel group or "none"
Access-List-Inbound		Y	Y	86	String	Single	Access list ID
Access-List-Outbound		Y	Y	87	String	Single	Access list ID
Perfect-Forward-Secrecy-Enable	Y	Y	Y	88	Boolean	Single	0 = No 1 = Yes
NAC-Enable	Y			89	Integer		0 = No 1 = Yes
NAC-Status-Query-Timer	Y			90	Integer		30 - 1800 seconds
NAC-Revalidation-Timer	Y			91	Integer		300 - 86400 seconds
NAC-Default-ACL	Y			92	String		Access list

Table E-4 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/Type	Single or Multi-Valued	Description or Value
WebVPN-URL-Entry-Enable	Y	Y		93	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Access-Enable	Y	Y		94	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Entry-Enable	Y	Y		95	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Browsing-Enable	Y	Y		96	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-Enable	Y	Y		97	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Outlook-Exchange-Proxy-Enable	Y	Y		98	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-HTTP-Proxy	Y	Y		99	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Auto-Applet-Download-Enable	Y	Y		100	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Citrix-Metaframe-Enable	Y	Y		101	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Apply-ACL	Y	Y		102	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Enable	Y	Y		103	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Required	Y	Y		104	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Keep-Installation	Y	Y		105	Integer	Single	0 = Disabled 1 = Enabled
Strip-Realm	Y	Y	Y	135	Boolean	Single	0 = Disabled 1 = Enabled

**Note**

RADIUS attribute names do not contain the cVPN3000 prefix to better reflect support for all three security appliances (VPN 3000, PIX, and the ASA). Cisco Secure ACS 4.x supports this new nomenclature, but attribute names in pre-4.0 ACS releases still include the cVPN3000 prefix. The appliances enforce the RADIUS attributes based on attribute numeric ID, not attribute name. LDAP attributes are enforced by their name, not by the ID.

Security Appliance TACACS+ Attributes

The security appliance provides support for TACACS+ attributes. TACACS+ separates the functions of authentication, authorization, and accounting. The protocol supports two types of attributes: mandatory and optional. Both the server and client must understand a mandatory attribute, and the mandatory attribute must be applied to the user. An optional attribute may or may not be understood or used.



Note

To use TACACS+ attributes, make sure you have enabled AAA services on the NAS.

[Table E-5](#) lists supported TACACS+ authorization response attributes for cut-through-proxy connections. [Table E-6](#) lists supported TACACS+ accounting attributes.

Table E-5 *Supported TACACS+ Authorization Response Attributes*

Attribute	Description
acl	Identifies a locally configured access list to be applied to the connection.
idletime	Indicates the amount of inactivity in minutes that is allowed before the authenticated user session is terminated.
timeout	Specifies the absolute amount of time in minutes that authentication credentials remain active before the authenticated user session is terminated.

Table E-6 *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's 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.



GLOSSARY

[Numerics](#) | [A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) | [X](#)

Numerics

3DES See [DES](#).

A

AAA Authentication, authorization, and accounting. See also [TACACS+](#) and [RADIUS](#).

ABR Area Border Router. In [OSPF](#), a router with interfaces in multiple areas.

ACE Access Control Entry. Information entered into the configuration that lets you specify what type of traffic to permit or deny on an [interface](#). By default, traffic that is not explicitly permitted is denied.

Access Modes The security appliance CLI uses several command modes. The commands available in each mode vary. See also [user EXEC mode](#), [privileged EXEC mode](#), [global configuration mode](#), [command-specific configuration mode](#).

ACL Access Control List. A collection of [ACEs](#). An ACL lets you specify what type of traffic to allow on an interface. By default, traffic that is not explicitly permitted is denied. ACLs are usually applied to the [interface](#) which is the source of inbound traffic. See also [rule](#), [outbound ACL](#).

ActiveX A set of object-oriented programming technologies and tools used to create mobile or portable programs. An ActiveX program is roughly equivalent to a Java applet.

Address Resolution Protocol See [ARP](#).

address translation The translation of a network address and/or port to another network address/or port. See also [IP address](#), [interface PAT](#), [NAT](#), [PAT](#), [Static PAT](#), [xlate](#).

AES Advanced Encryption Standard. A symmetric block cipher that can encrypt and decrypt information. The AES algorithm is capable of using cryptographic keys of 128, 192 and 256 bits to encrypt and decrypt data in blocks of 128 bits. See also [DES](#).

AH Authentication Header. An IP protocol (type 51) that can ensure data integrity, authentication, and replay detection. AH is embedded in the data to be protected (a full IP datagram, for example). AH can be used either by itself or with [ESP](#). This is an older [IPSec](#) protocol that is less important in most networks than [ESP](#). AH provides authentication services but does not provide encryption services. It is provided to ensure compatibility with [IPSec](#) peers that do not support [ESP](#), which provides both [authentication](#) and [encryption](#). See also [encryption](#) and [VPN](#). Refer to the RFC 2402.

A record address “A” stands for address, and refers to name-to-address mapped records in [DNS](#).

APCF	Application Profile Customization Framework. Lets the security appliance handle non-standard applications so that they render correctly over a WebVPN connection.
ARP	Address Resolution Protocol. A low-level TCP/IP protocol that maps a hardware address, or MAC address, to an IP address. An example hardware address is 00:00:a6:00:01:ba. The first three groups of characters (00:00:a6) identify the manufacturer; the rest of the characters (00:01:ba) identify the system card. ARP is defined in RFC 826.
ASA	Adaptive Security Algorithm. Used by the security appliance to perform inspections. ASA allows one-way (inside to outside) connections without an explicit configuration for each internal system and application. See also inspection engine .
ASA	adaptive security appliance.
ASDM	Adaptive Security Device Manager. An application for managing and configuring a single security appliance.
asymmetric encryption	Also called public key systems, asymmetric encryption allows anyone to obtain access to the public key of anyone else. Once the public key is accessed, one can send an encrypted message to that person using the public key. See also encryption , public key .
authentication	Cryptographic protocols and services that verify the identity of users and the integrity of data. One of the functions of the IPSec framework. Authentication establishes the integrity of datastream and ensures that it is not tampered with in transit. It also provides confirmation about the origin of the datastream. See also AAA , encryption , and VPN .
Auto Applet Download	Automatically downloads the WebVPN port-forwarding applet when the user first logs in to WebVPN.
auto-signon	This command provides a single sign-on method for WebVPN users. It passes the WebVPN login credentials (username and password) to internal servers for authentication using NTLM authentication, basic authentication, or both.

B

Backup Server	IPSec backup servers let a VPN client connect to the central site when the primary security appliance is unavailable.
BGP	Border Gateway Protocol. BGP performs interdomain routing in TCP/IP networks. BGP is an Exterior Gateway Protocol, which means that it performs routing between multiple autonomous systems or domains and exchanges routing and access information with other BGP systems. The security appliance does not support BGP. See also EGP .
BLT stream	Bandwidth Limited Traffic stream. Stream or flow of packets whose bandwidth is constrained.
BOOTP	Bootstrap Protocol. Lets diskless workstations boot over the network as is described in RFC 951 and RFC 1542.
BPDU	Bridge Protocol Data Unit. Spanning-Tree Protocol hello packet that is sent out at configurable intervals to exchange information among bridges in the network. Protocol data unit is the OSI term for packet.

C

CA	Certificate Authority, Certification Authority. A third-party entity that is responsible for issuing and revoking certificates. Each device with the public key of the CA can authenticate a device that has a certificate issued by the CA. The term CA also refers to software that provides CA services. See also certificate , CRL , public key , RA .
cache	A temporary repository of information accumulated from previous task executions that can be reused, decreasing the time required to perform the tasks. Caching stores frequently reused objects in the system cache, which reduces the need to perform repeated rewriting and compressing of content.
CBC	Cipher Block Chaining. A cryptographic technique that increases the encryption strength of an algorithm. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPSec packet.
certificate	A signed cryptographic object that contains the identity of a user or device and the public key of the CA that issued the certificate. Certificates have an expiration date and may also be placed on a CRL if known to be compromised. Certificates also establish non-repudiation for IKE negotiation, which means that you can prove to a third party that IKE negotiation was completed with a specific peer.
CHAP	Challenge Handshake Authentication Protocol.
CIFS	Common Internet File System. It is a platform-independent file sharing system that provides users with network access to files, printers, and other machine resources. Microsoft implemented CIFS for networks of Windows computers, however, open source implementations of CIFS provide file access to servers running other operating systems, such as Linux, UNIX, and Mac OS X.
Citrix	An application that virtualizes client-server applications and optimizes web applications.
CLI	command line interface. The primary interface for entering configuration and monitoring commands to the security appliance.
client/server computing	Distributed computing (processing) network systems in which transaction responsibilities are divided into two parts: client (front end) and server (back end). Also called distributed computing. See also RPC .
Client update	Lets you update revisions of clients to which the update applies; provide a URL or IP address from which to get the update; and, in the case of Windows clients, optionally notify users that they should update their VPN client version.
command-specific configuration mode	From global configuration mode, some commands enter a command-specific configuration mode. All user EXEC, privileged EXEC, global configuration, and command-specific configuration commands are available in this mode. See also global configuration mode , privileged EXEC mode , user EXEC mode .
Compression	The process of encoding information using fewer bits or other information-bearing units than an unencoded representation would use. Compression can reduce the size of transferring packets and increase communication performance.
configuration, config, config file	A file on the security appliance that represents the equivalent of settings, preferences, and properties administered by ASDM or the CLI .

Content Rewriting/Transformation	Interprets and modifies applications so that they render correctly over a WebVPN connection.
cookie	A cookie is a object stored by a browser. Cookies contain information, such as user preferences, to persistent storage.
CPU	Central Processing Unit. Main processor.
CRC	Cyclical Redundancy Check. Error-checking technique in which the frame recipient calculates a remainder by dividing frame contents by a prime binary divisor and compares the calculated remainder to a value stored in the frame by the sending node.
CRL	Certificate Revocation List. A digitally signed message that lists all of the current but revoked certificates listed by a given CA . This is analogous to a book of stolen charge card numbers that allow stores to reject bad credit cards. When certificates are revoked, they are added to a CRL. When you implement authentication using certificates, you can choose to use CRLs or not. Using CRLs lets you easily revoke certificates before they expire, but the CRL is generally only maintained by the CA or an RA . If you are using CRLs and the connection to the CA or RA is not available when authentication is requested, the authentication request will fail. See also CA , certificate , public key , RA .
CRV	Call Reference Value. Used by H.225.0 to distinguish call legs signalled between two entities.
cryptography	Encryption, authentication, integrity, keys and other services used for secure communication over networks. See also VPN and IPSec .
crypto map	A data structure with a unique name and sequence number that is used for configuring VPNs on the security appliance. A crypto map selects data flows that need security processing and defines the policy for these flows and the crypto peer that traffic needs to go to. A crypto map is applied to an interface. Crypto maps contain the ACLs , encryption standards, peers, and other parameters necessary to specify security policies for VPNs using IKE and IPSec . See also VPN .
CTIQBE	Computer Telephony Interface Quick Buffer Encoding. A protocol used in IP telephony between the Cisco CallManager and CTI TAPI and JTAPI applications. CTIQBE is used by the TAPI/JTAPI protocol inspection module and supports NAT , PAT , and bi-directional NAT . This enables Cisco IP SoftPhone and other Cisco TAPI/JTAPI applications to communicate with Cisco CallManager for call setup and voice traffic across the security appliance.
cut-through proxy	Enables the security appliance to provide faster traffic flow after user authentication. The cut-through proxy challenges a user initially at the application layer. After the security appliance authenticates the user, it shifts the session flow and all traffic flows directly and quickly between the source and destination while maintaining session state information.

D

data confidentiality	Describes any method that manipulates data so that no attacker can read it. This is commonly achieved by data encryption and keys that are only available to the parties involved in the communication.
data integrity	Describes mechanisms that, through the use of encryption based on secret key or public key algorithms, allow the recipient of a piece of protected data to verify that the data has not been modified in transit.

data origin authentication	A security service where the receiver can verify that protected data could have originated only from the sender. This service requires a data integrity service plus a key distribution mechanism, where a secret key is shared only between the sender and receiver.
decryption	Application of a specific algorithm or cipher to encrypted data so as to render the data comprehensible to those who are authorized to see the information. See also encryption .
DES	Data encryption standard. DES was published in 1977 by the National Bureau of Standards and is a secret key encryption scheme based on the Lucifer algorithm from IBM. Cisco uses DES in classic crypto (40-bit and 56-bit key lengths), IPSec crypto (56-bit key), and 3DES (triple DES), which performs encryption three times using a 56-bit key. 3DES is more secure than DES but requires more processing for encryption and decryption. See also AES , ESP .
DHCP	Dynamic Host Configuration Protocol. Provides a mechanism for allocating IP addresses to hosts dynamically, so that addresses can be reused when hosts no longer need them and so that mobile computers, such as laptops, receive an IP address applicable to the LAN to which it is connected.
Diffie-Hellman	A public key cryptography protocol that allows two parties to establish a shared secret over insecure communications channels. Diffie-Hellman is used within IKE to establish session keys. Diffie-Hellman is a component of Oakley key exchange.
Diffie-Hellman Group 1, Group 2, Group 5, Group 7	Diffie-Hellman refers to a type of public key cryptography using asymmetric encryption based on large prime numbers to establish both Phase 1 and Phase 2 SAs . Group 1 provides a smaller prime number than Group 2 but may be the only version supported by some IPSec peers. Diffie-Hellman Group 5 uses a 1536-bit prime number, is the most secure, and is recommended for use with AES . Group 7 has an elliptical curve field size of 163 bits and is for use with the Movian VPN client, but works with any peer that supports Group 7 (ECC). See also VPN and encryption .
digital certificate	See certificate .
DMZ	See interface .
DN	Distinguished Name. Global, authoritative name of an entry in the OSI Directory (X.500).
DNS	Domain Name System (or Service). An Internet service that translates domain names into IP addresses.
DoS	Denial of Service. A type of network attack in which the goal is to render a network service unavailable.
DSL	digital subscriber line. Public network technology that delivers high bandwidth over conventional copper wiring at limited distances. DSL is provisioned via modem pairs, with one modem located at a central office and the other at the customer site. Because most DSL technologies do not use the whole bandwidth of the twisted pair, there is room remaining for a voice channel.
DSP	digital signal processor. A DSP segments a voice signal into frames and stores them in voice packets.
DSS	Digital Signature Standard. A digital signature algorithm designed by The US National Institute of Standards and Technology and based on public-key cryptography. DSS does not do user datagram encryption. DSS is a component in classic crypto, as well as the Redcreek IPSec card, but not in IPSec implemented in Cisco IOS software.

Dynamic NAT	See NAT and address translation .
Dynamic PAT	Dynamic Port Address Translation. Dynamic PAT lets multiple outbound sessions appear to originate from a single IP address. With PAT enabled, the security appliance chooses a unique port number from the PAT IP address for each outbound translation slot (xlate). This feature is valuable when an ISP cannot allocate enough unique IP addresses for your outbound connections. The global pool addresses always come first, before a PAT address is used. See also NAT , Static PAT , and xlate .
E	
ECHO	See Ping , ICMP . See also inspection engine .
EGP	Exterior Gateway Protocol. Replaced by BGP. The security appliance does not support EGP. See also BGP .
EIGRP	Enhanced Interior Gateway Routing Protocol. The security appliance does not support EIGRP.
EMBLEM	Enterprise Management BaseLine Embedded Manageability. A syslog format designed to be consistent with the Cisco IOS system log format and is more compatible with CiscoWorks management applications.
encryption	Application of a specific algorithm or cipher to data so as to render the data incomprehensible to those unauthorized to see the information. See also decryption .
ESMTP	Extended SMTP . Extended version of SMTP that includes additional functionality, such as delivery notification and session delivery. ESMTP is described in RFC 1869, SMTP Service Extensions.
ESP	Encapsulating Security Payload. An IPSec protocol, ESP provides authentication and encryption services for establishing a secure tunnel over an insecure network. For more information, refer to RFCs 2406 and 1827.
F	
failover, failover mode	Failover lets you configure two security appliances so that one will take over operation if the other one fails. The security appliance supports two failover configurations, Active/Active failover and Active/Standby failover. Each failover configuration has its own method for determining and performing failover. With Active/Active failover, both units can pass network traffic. This lets you configure load balancing on your network. Active/Active failover is only available on units running in multiple context mode. With Active/Standby failover, only one unit passes traffic while the other unit waits in a standby state. Active/Standby failover is available on units running in either single or multiple context mode.
Fixup	See inspection engine .
Flash, Flash memory	A nonvolatile storage device used to store the configuration file when the security appliance is powered down.
FQDN/IP	Fully qualified domain name/IP address. IPSec parameter that identifies peers that are security gateways.

FragGuard Provides IP fragment protection and performs full reassembly of all [ICMP](#) error messages and virtual reassembly of the remaining IP fragments that are routed through the security appliance.

FTP File Transfer Protocol. Part of the TCP/IP protocol stack, used for transferring files between hosts.

G

GGSN gateway [GPRS](#) support node. A wireless gateway that allows mobile cell phone users to access the public data network or specified private IP networks.

global configuration mode Global configuration mode lets you to change the security appliance configuration. All user EXEC, privileged EXEC, and global configuration commands are available in this mode. See also [user EXEC mode](#), [privileged EXEC mode](#), [command-specific configuration mode](#).

GMT Greenwich Mean Time. Replaced by UTC (Coordinated Universal Time) in 1967 as the world time standard.

GPRS general packet radio service. A service defined and standardized by the European Telecommunication Standards Institute. GPRS is an IP-packet-based extension of [GSM](#) networks and provides mobile, wireless, data communications

GRE Generic Routing Encapsulation described in RFCs 1701 and 1702. GRE is a tunneling protocol that can encapsulate a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link to routers at remote points over an IP network. By connecting multiprotocol subnetworks in a single-protocol backbone environment, IP tunneling using GRE allows network expansion across a single protocol backbone environment.

GSM Global System for Mobile Communication. A digital, mobile, radio standard developed for mobile, wireless, voice communications.

GTP GPRS tunneling protocol. GTP handles the flow of user packet data and signaling information between the [SGSN](#) and [GGSN](#) in a [GPRS](#) network. GTP is defined on both the Gn and Gp interfaces of a [GPRS](#) network.

H

H.225 A protocol used for TCP signalling in applications such as video conferencing. See also [H.323](#) and [inspection engine](#).

H.225.0 An ITU standard that governs H.225.0 session establishment and packetization. H.225.0 actually describes several different protocols: RAS, use of Q.931, and use of [RTP](#).

H.245 An ITU standard that governs H.245 endpoint control.

H.320 Suite of ITU-T standard specifications for video conferencing over circuit-switched media, such as ISDN, fractional T-1, and switched-56 lines. Extensions of ITU-T standard H.320 enable video conferencing over LANs and other packet-switched networks, as well as video over the [Internet](#).

H.323	Allows dissimilar communication devices to communicate with each other by using a standardized communication protocol. H.323 defines a common set of CODECs, call setup and negotiating procedures, and basic data transport methods.
H.323 RAS	Registration, admission, and status signaling protocol. Enables devices to perform registration, admissions, bandwidth changes, and status and disengage procedures between VoIP gateway and the gatekeeper.
H.450.2	Call transfer supplementary service for H.323 .
H.450.3	Call diversion supplementary service for H.323 .
Hash, Hash Algorithm	A hash algorithm is a one way function that operates on a message of arbitrary length to create a fixed-length message digest used by cryptographic services to ensure its data integrity. MD5 has a smaller digest and is considered to be slightly faster than SHA-1 . Cisco uses both SHA-1 and MD5 hashes within our implementation of the IPSec framework. See also encryption , HMAC , and VPN .
headend	A firewall, concentrator, or other host that serves as the entry point into a private network for VPN client connections over the public network. See also ISP and VPN .
HMAC	A mechanism for message authentication using cryptographic hashes such as SHA-1 and MD5 .
host	The name for any device on a TCP/IP network that has an IP address. See also network and node .
host/network	An IP address and netmask used with other information to identify a single host or network subnet for security appliance configuration, such as an address translation (xlate) or ACE .
HTTP	Hypertext Transfer Protocol. A protocol used by browsers and web servers to transfer files. When a user views a web page, the browser can use HTTP to request and receive the files used by the web page. HTTP transmissions are not encrypted.
HTTPS	Hypertext Transfer Protocol Secure. An SSL -encrypted version of HTTP.

IANA	Internet Assigned Number Authority. Assigns all port and protocol numbers for use on the Internet .
ICMP	Internet Control Message Protocol. Network-layer Internet protocol that reports errors and provides other information relevant to IP packet processing.
IDS	Intrusion Detection System. A method of detecting malicious network activity by signatures and then implementing a policy for that signature.
IETF	The Internet Engineering Task Force. A technical standards organization that develops RFC documents defining protocols for the Internet .
IGMP	Internet Group Management Protocol. IGMP is a protocol used by IPv4 systems to report IP multicast memberships to neighboring multicast routers.

IKE	Internet Key Exchange. IKE establishes a shared security policy and authenticates keys for services (such as IPSec) that require keys. Before any IPSec traffic can be passed, each security appliance must verify the identity of its peer. This can be done by manually entering preshared keys into both hosts or by a CA service. IKE is a hybrid protocol that uses part Oakley and part of another protocol suite called SKEME inside ISAKMP framework. This is the protocol formerly known as ISAKMP/Oakley, and is defined in RFC 2409.
IKE Extended Authentication	IKE Extended Authenticate (Xauth) is implemented per the IETF draft-ietf-ipsec-isakmp-xauth-04.txt (“extended authentication” draft). This protocol provides the capability of authenticating a user within IKE using TACACS+ or RADIUS .
IKE Mode Configuration	IKE Mode Configuration is implemented per the IETF draft-ietf-ipsec-isakmp-mode-cfg-04.txt. IKE Mode Configuration provides a method for a security gateway to download an IP address (and other network level configuration) to the VPN client as part of an IKE negotiation.
ILS	Internet Locator Service. ILS is based on LDAP and is ILSv2 compliant. ILS was developed by Microsoft for use with its NetMeeting, SiteServer, and Active Directory products.
IMAP	Internet Message Access Protocol. Method of accessing e-mail or bulletin board messages kept on a mail server that can be shared. IMAP permits client e-mail applications to access remote message stores as if they were local without actually transferring the message.
implicit rule	An access rule automatically created by the security appliance based on default rules or as a result of user-defined rules.
IMSI	International Mobile Subscriber Identity. One of two components of a GTP tunnel ID, the other being the NSAPI . See also NSAPI .
inside	The first interface, usually port 1, that connects your internal, “trusted” network protected by the security appliance. See also interface , interface names .
inspection engine	The security appliance inspects certain application-level protocols to identify the location of embedded addressing information in traffic. This allows NAT to translate these embedded addresses and to update any checksum or other fields that are affected by the translation. Because many protocols open secondary TCP or UDP ports, each application inspection engine also monitors sessions to determine the port numbers for secondary channels. The initial session on a well-known port is used to negotiate dynamically assigned port numbers. The application inspection engine monitors these sessions, identifies the dynamic port assignments, and permits data exchange on these ports for the duration of the specific session. Some of the protocols that the security appliance can inspect are CTIQBE , FTP , H.323 , HTTP , MGCP , SMTP , and SNMP .
interface	The physical connection between a particular network and a security appliance.
interface ip_address	The IP address of a security appliance network interface. Each interface IP address must be unique. Two or more interfaces must not be given the same IP address or IP addresses that are on the same IP network.
interface names	Human readable name assigned to a security appliance network interface. The inside interface default name is “inside” and the outside interface default name is “outside.” Any perimeter interface default names are “intf <i>n</i> ”, such as intf2 for the first perimeter interface, intf3 for the second perimeter interface, and so on to the last interface. The numbers in the intf string corresponds to the position of the interface card in the security appliance. You can use the default names or, if you are an experienced user, give each interface a more meaningful name. See also inside , intf<i>n</i> , outside .

intfn	Any interface, usually beginning with port 2, that connects to a subset network of your design that you can custom name and configure.
interface PAT	The use of PAT where the PAT IP address is also the IP address of the outside interface. See Dynamic PAT , Static PAT .
Internet	The global network that uses IP . Not a LAN . See also intranet .
intranet	Intranetwork. A LAN that uses IP . See also network and Internet .
IP	Internet Protocol. IP protocols are the most popular nonproprietary protocols because they can be used to communicate across any set of interconnected networks and are equally well suited for LAN and WAN communications.
IPS	Intrusion Prevention Service. An in-line, deep-packet inspection-based solution that helps mitigate a wide range of network attacks.
IP address	An IP protocol address. A security appliance interface <code>ip_address</code> . IP version 4 addresses are 32 bits in length. This address space is used to designate the network number, optional subnetwork number, and a host number. The 32 bits are grouped into four octets (8 binary bits), represented by 4 decimal numbers separated by periods, or dots. The meaning of each of the four octets is determined by their use in a particular network.
IP pool	A range of local IP addresses specified by a name, and a range with a starting IP address and an ending address. IP Pools are used by DHCP and VPNs to assign local IP addresses to clients on the inside interface.
IPSec	IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPSec provides these security services at the IP layer. IPSec uses IKE to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPSec. IPSec can protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.
IPSec Phase 1	The first phase of negotiating IPSec , includes the key exchange and the ISAKMP portions of IPSec .
IPSec Phase 2	The second phase of negotiating IPSec . Phase two determines the type of encryption rules used for payload, the source and destination that will be used for encryption, the definition of interesting traffic according to access lists, and the IPSec peer. IPSec is applied to the interface in Phase 2.
IPSec transform set	A transform set specifies the IPSec protocol, encryption algorithm, and hash algorithm to use on traffic matching the IPSec policy. A transform describes a security protocol (AH or ESP) with its corresponding algorithms. The IPSec protocol used in almost all transform sets is ESP with the DES algorithm and HMAC-SHA for authentication.
ISAKMP	Internet Security Association and Key Management Protocol. A protocol framework that defines payload formats, the mechanics of implementing a key exchange protocol, and the negotiation of a security association. See IKE .
ISP	Internet Service Provider. An organization that provides connection to the Internet via their services, such as modem dial in over telephone voice lines or DSL .

J

JTAPI Java Telephony Application Programming Interface. A Java-based API supporting telephony functions. See also [TAPI](#).

K

key A data object used for [encryption](#), [decryption](#), or [authentication](#).

L

LAN Local area network. A network residing in one location, such as a single building or campus. See also [Internet](#), [intranet](#), and [network](#).

layer, layers Networking models implement layers with which different protocols are associated. The most common networking model is the OSI model, which consists of the following 7 layers, in order: physical, data link, network, transport, session, presentation, and application.

LCN Logical channel number.

LDAP Lightweight Directory Access Protocol. LDAP provides management and browser applications with access to X.500 directories.

M

mask A 32-bit mask that shows how an [Internet](#) address is divided into network, subnet, and host parts. The mask has ones in the bit positions to be used for the network and subnet parts, and zeros for the host part. The mask should contain at least the standard network portion, and the subnet field should be contiguous with the network portion.

MCR See [multicast](#).

MC router Multicast (MC) routers route multicast data transmissions to the hosts on each LAN in an internetwork that are registered to receive specific multimedia or other broadcasts. See also [multicast](#).

MD5 Message Digest 5. A one-way hashing algorithm that produces a 128-bit hash. Both MD5 and [SHA-1](#) are variations on MD4 and are designed to strengthen the security of the MD4 hashing algorithm. [SHA-1](#) is more secure than MD4 and MD5. Cisco uses hashes for authentication within the [IPSec](#) framework. Also used for message authentication in SNMP v.2. MD5 verifies the integrity of the communication, authenticates the origin, and checks for timeliness. [MD5](#) has a smaller digest and is considered to be slightly faster than [SHA-1](#).

MDI Media dependent interface.

MDIX Media dependent interface crossover.

Message Digest	A message digest is created by a hash algorithm, such as MD5 or SHA-1 , that is used for ensuring message integrity.
MGCP	Media Gateway Control Protocol. Media Gateway Control Protocol is a protocol for the control of VoIP calls by external call-control elements known as media gateway controllers or call agents. MGCP merges the IPDC and SGCP protocols.
Mode	See Access Modes .
Mode Config	See IKE Mode Configuration .
Modular Policy Framework	Modular Policy Framework. A means of configuring security appliance features in a manner to similar to Cisco IOS software Modular QoS CLI.
MS	mobile station. Refers generically to any mobile device, such as a mobile handset or computer, that is used to access network services. GPRS networks support three classes of MS, which describe the type of operation supported within the GPRS and the GSM mobile wireless networks. For example, a Class A MS supports simultaneous operation of GPRS and GSM services.
MS-CHAP	Microsoft CHAP .
MTU	Maximum transmission unit, the maximum number of bytes in a packet that can flow efficiently across the network with best response time. For Ethernet, the default MTU is 1500 bytes, but each network can have different values, with serial connections having the smallest values. The MTU is described in RFC 1191.
multicast	Multicast refers to a network addressing method in which the source transmits a packet to multiple destinations, a multicast group, simultaneously. See also PIM , SMR .

N	
N2H2	A third-party, policy-oriented filtering application that works with the security appliance to control user web access. N2H2 can filter HTTP requests based on destination host name, destination IP address, and username and password. The N2H2 corporation was acquired by Secure Computing in October, 2003.
NAT	Network Address Translation. Mechanism for reducing the need for globally unique IP addresses. NAT allows an organization with addresses that are not globally unique to connect to the Internet by translating those addresses into a globally routable address space.
NEM	Network Extension Mode. Lets VPN hardware clients present a single, routable network to the remote private network over the VPN tunnel.
NetBIOS	Network Basic Input/Output System. A Microsoft protocol that supports Windows host name registration, session management, and data transfer. The security appliance supports NetBIOS by performing NAT of the packets for NBNS UDP port 137 and NBDS UDP port 138.
netmask	See mask .
network	In the context of security appliance configuration, a network is a group of computing devices that share part of an IP address space and not a single host. A network consists of multiple nodes or hosts. See also host , Internet , intranet , IP , LAN , and node .

NMS	network management system. System responsible for managing at least part of a network. An NMS is generally a reasonably powerful and well-equipped computer, such as an engineering workstation. NMSs communicate with agents to help keep track of network statistics and resources.
node	Devices such as routers and printers that would not normally be called hosts. See also host , network .
nonvolatile storage, memory	Storage or memory that, unlike RAM, retains its contents without power. Data in a nonvolatile storage device survives a power-off, power-on cycle or reboot.
NSAPI	Network service access point identifier. One of two components of a GTP tunnel ID, the other component being the IMSI . See also IMSI .
NSSA	Not-so-stubby-area. An OSPF feature described by RFC 1587. NSSA was first introduced in Cisco IOS software release 11.2. It is a non-proprietary extension of the existing stub area feature that allows the injection of external routes in a limited fashion into the stub area.
NTLM	NT Lan Manager. A Microsoft Windows challenge-response authentication method.
NTP	Network time protocol.

O

Oakley	A key exchange protocol that defines how to acquire authenticated keying material. The basic mechanism for Oakley is the Diffie-Hellman key exchange algorithm. Oakley is defined in RFC 2412.
object grouping	Simplifies access control by letting you apply access control statements to groups of network objects, such as protocol, services, hosts, and networks.
OSPF	Open Shortest Path First. OSPF is a routing protocol for IP networks. OSPF is a routing protocol widely deployed in large networks because of its efficient use of network bandwidth and its rapid convergence after changes in topology. The security appliance supports OSPF.
OU	Organizational Unit. An X.500 directory attribute.
outbound	Refers to traffic whose destination is on an interface with lower security than the source interface.
outbound ACL	An ACL applied to outbound traffic.
outside	The first interface, usually port 0, that connects to other “untrusted” networks outside the security appliance; the Internet . See also interface , interface names , outbound .

P

PAC	PPTP Access Concentrator. A device attached to one or more PSTN or ISDN lines capable of PPP operation and of handling the PPTP protocol. The PAC need only implement TCP/IP to pass traffic to one or more PNSs . It may also tunnel non-IP protocols.
PAT	See Dynamic PAT , interface PAT , and Static PAT .
PDP	Packet Data Protocol.

Perfmon	The security appliance feature that gathers and reports a wide variety of feature statistics, such as connections/second, xlates/second, etc.
PFS	Perfect Forwarding Secrecy. PFS enhances security by using different security key for the IPSec Phase 1 and Phase 2 SAs . Without PFS, the same security key is used to establish SAs in both phases. PFS ensures that a given IPSec SA key was not derived from any other secret (like some other keys). In other words, if someone were to break a key, PFS ensures that the attacker would not be able to derive any other key. If PFS were not enabled, someone could hypothetically break the IKE SA secret key, copy all the IPSec protected data, and then use knowledge of the IKE SA secret to compromise the IPSec SA setup by this IKE SA . With PFS, breaking IKE would not give an attacker immediate access to IPSec . The attacker would have to break each IPSec SA individually.
Phase 1	See IPSec Phase 1 .
Phase 2	See IPSec Phase 2 .
PIM	Protocol Independent Multicast. PIM provides a scalable method for determining the best paths for distributing a specific multicast transmission to a group of hosts. Each host has registered using IGMP to receive the transmission. See also PIM-SM .
PIM-SM	Protocol Independent Multicast-Sparse Mode. With PIM-SM, which is the default for Cisco routers, when the source of a multicast transmission begins broadcasting, the traffic is forwarded from one MC router to the next, until the packets reach every registered host. See also PIM .
Ping	An ICMP request sent by a host to determine if a second host is accessible.
PIX	Private Internet eXchange. The Cisco PIX 500-series security appliances range from compact, plug-and-play desktop models for small/home offices to carrier-class gigabit models for the most demanding enterprise and service provider environments. Cisco PIX security appliances provide robust, enterprise-class integrated network security services to create a strong multilayered defense for fast changing network environments.
PKCS12	A standard for the transfer of PKI-related data, such as private keys, certificates, and other data. Devices supporting this standard let administrators maintain a single set of personal identity information.
PNS	PPTP Network Server. A PNS is envisioned to operate on general-purpose computing/server platforms. The PNS handles the server side of PPTP . Because PPTP relies completely on TCP/IP and is independent of the interface hardware, the PNS may use any combination of IP interface hardware including LAN and WAN devices.
Policy NAT	Lets you identify local traffic for address translation by specifying the source and destination addresses (or ports) in an access list.
POP	Post Office Protocol. Protocol that client e-mail applications use to retrieve mail from a mail server.
Pool	See IP pool .
Port	A field in the packet headers of TCP and UDP protocols that identifies the higher level service which is the source or destination of the packet.
PPP	Point-to-Point Protocol. Developed for dial-up ISP access using analog phone lines and modems.

PPTP	Point-to-Point Tunneling Protocol. PPTP was introduced by Microsoft to provide secure remote access to Windows networks; however, because it is vulnerable to attack, PPTP is commonly used only when stronger security methods are not available or are not required. PPTP Ports are pptp, 1723/tcp, 1723/udp, and pptp. For more information about PPTP, see RFC 2637. See also PAC , PPTP GRE , PPTP GRE tunnel , PNS , PPTP session , and PPTP TCP .
PPTP GRE	Version 1 of GRE for encapsulating PPP traffic.
PPTP GRE tunnel	A tunnel defined by a PNS-PAC pair. The tunnel protocol is defined by a modified version of GRE . The tunnel carries PPP datagrams between the PAC and the PNS . Many sessions are multiplexed on a single tunnel. A control connection operating over TCP controls the establishment, release, and maintenance of sessions and of the tunnel itself.
PPTP session	PPTP is connection-oriented. The PNS and PAC maintain state for each user that is attached to a PAC . A session is created when end-to-end PPP connection is attempted between a dial user and the PNS . The datagrams related to a session are sent over the tunnel between the PAC and PNS .
PPTP TCP	Standard TCP session over which PPTP call control and management information is passed. The control session is logically associated with, but separate from, the sessions being tunneled through a PPTP tunnel.
preshared key	A preshared key provides a method of IKE authentication that is suitable for networks with a limited, static number of IPSec peers. This method is limited in scalability because the key must be configured for each pair of IPSec peers. When a new IPSec peer is added to the network, the preshared key must be configured for every IPSec peer with which it communicates. Using certificates and CAs provides a more scalable method of IKE authentication.
primary, primary unit	The security appliance normally operating when two units, a primary and secondary, are operating in failover mode.
privileged EXEC mode	Privileged EXEC mode lets you to change current settings. Any user EXEC mode command will work in privileged EXEC mode. See also command-specific configuration mode , global configuration mode , user EXEC mode .
protocol, protocol literals	A standard that defines the exchange of packets between network nodes for communication. Protocols work together in layers. Protocols are specified in a security appliance configuration as part of defining a security policy by their literal values or port numbers. Possible security appliance protocol literal values are ahp, eigrp, esp, gre, icmp, igmp, igmp, ip, ipinip, ipsec, nos, ospf, pcp, snp, tcp, and udp.
Proxy-ARP	Enables the security appliance to reply to an ARP request for IP addresses in the global pool. See also ARP .
public key	A public key is one of a pair of keys that are generated by devices involved in public key infrastructure. Data encrypted with a public key can only be decrypted using the associated private key. When a private key is used to produce a digital signature, the receiver can use the public key of the sender to verify that the message was signed by the sender. These characteristics of key pairs provide a scalable and secure method of authentication over an insecure media, such as the Internet .

Q

QoS quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability.

R

RA Registration Authority. An authorized proxy for a [CA](#). RAs can perform certificate enrollment and can issue [CRLs](#). See also [CA](#), [certificate](#), [public key](#).

RADIUS Remote Authentication Dial-In User Service. RADIUS is a distributed client/server system that secures networks against unauthorized access. RFC 2058 and RFC 2059 define the RADIUS protocol standard. See also [AAA](#) and [TACACS+](#).

Refresh Retrieve the running configuration from the security appliance and update the screen. The icon and the button perform the same function.

registration authority See [RA](#).

replay-detection A security service where the receiver can reject old or duplicate packets to defeat replay attacks. Replay attacks rely on the attacker sending out older or duplicate packets to the receiver and the receiver thinking that the bogus traffic is legitimate. Replay-detection is done by using sequence numbers combined with authentication, and is a standard feature of [IPSec](#).

RFC Request for Comments. RFC documents define protocols and standards for communications over the [Internet](#). RFCs are developed and published by [IETF](#).

RIP Routing Information Protocol. Interior gateway protocol (IGP) supplied with UNIX BSD systems. The most common IGP in the [Internet](#). RIP uses hop count as a routing metric.

RLLA Reserved Link Local Address. Multicast addresses range from 224.0.0.0 to 239.255.255.255, however only the range 224.0.1.0 to 239.255.255.255 is available to us. The first part of the multicast address range, 224.0.0.0 to 224.0.0.255, is reserved and referred to as the RLLA. These addresses are unavailable. We can exclude the RLLA range by specifying: 224.0.1.0 to 239.255.255.255. 224.0.0.0 to 239.255.255.255 excluding 224.0.0.0 to 224.0.0.255. This is the same as specifying: 224.0.1.0 to 239.255.255.255.

route, routing The path through a [network](#).

routed firewall mode In routed firewall mode, the security appliance is counted as a router hop in the network. It performs [NAT](#) between connected networks and can use [OSPF](#) or [RIP](#). See also [transparent firewall mode](#).

RPC Remote Procedure Call. RPCs are procedure calls that are built or specified by clients and executed on servers, with the results returned over the network to the clients.

RSA	A public key cryptographic algorithm (named after its inventors, Rivest, Shamir, and Adelman) with a variable key length. The main weakness of RSA is that it is significantly slow to compute compared to popular secret-key algorithms, such as DES . The Cisco implementation of IKE uses a Diffie-Hellman exchange to get the secret keys. This exchange can be authenticated with RSA (or preshared keys). With the Diffie-Hellman exchange, the DES key never crosses the network (not even in encrypted form), which is not the case with the RSA encrypt and sign technique. RSA is not public domain, and must be licensed from RSA Data Security.
RSH	Remote Shell. A protocol that allows a user to execute commands on a remote system without having to log in to the system. For example, RSH can be used to remotely examine the status of a number of access servers without connecting to each communication server, executing the command, and then disconnecting from the communication server.
RTCP	RTP Control Protocol. Protocol that monitors the QoS of an IPv6 RTP connection and conveys information about the on-going session. See also RTP .
RTP	Real-Time Transport Protocol. Commonly used with IP networks. RTP is designed to provide end-to-end network transport functions for applications transmitting real-time data, such as audio, video, or simulation data, over multicast or unicast network services. RTP provides such services as payload type identification, sequence numbering, timestamping, and delivery monitoring to real-time applications.
RTSP	Real Time Streaming Protocol. Enables the controlled delivery of real-time data, such as audio and video. RTSP is designed to work with established protocols, such as RTP and HTTP .
rule	Conditional statements added to the security appliance configuration to define security policy for a particular situation. See also ACE , ACL , NAT .
running configuration	The configuration currently running in RAM on the security appliance. The configuration that determines the operational characteristics of the security appliance.

S

SA	security association. An instance of security policy and keying material applied to a data flow. SAs are established in pairs by IPSec peers during both phases of IPSec . SAs specify the encryption algorithms and other security parameters used to create a secure tunnel. Phase 1 SAs (IKE SAs) establish a secure tunnel for negotiating Phase 2 SAs. Phase 2 SAs (IPSec SAs) establish the secure tunnel used for sending user data. Both IKE and IPSec use SAs, although SAs are independent of one another. IPSec SAs are unidirectional and they are unique in each security protocol. A set of SAs are needed for a protected data pipe, one per direction per protocol. For example, if you have a pipe that supports ESP between peers, one ESP SA is required for each direction. SAs are uniquely identified by destination (IPSec endpoint) address, security protocol (AH or ESP), and Security Parameter Index. IKE negotiates and establishes SAs on behalf of IPSec . A user can also establish IPSec SAs manually. An IKE SA is used by IKE only, and unlike the IPSec SA, it is bidirectional.
SCCP	Skinny Client Control Protocol. A Cisco-proprietary protocol used between Cisco Call Manager and Cisco VoIP phones.
SCEP	Simple Certificate Enrollment Protocol. A method of requesting and receiving (also known as enrolling) certificates from CAs .

SDP	Session Definition Protocol. An IETF protocol for the definition of Multimedia Services. SDP messages can be part of SGCP and MGCP messages.
secondary unit	The backup security appliance when two are operating in failover mode.
secret key	A secret key is a key shared only between the sender and receiver. See key , public key .
security context	You can partition a single security appliance into multiple virtual firewalls, known as security contexts. Each context is an independent firewall, with its own security policy, interfaces, and administrators. Multiple contexts are similar to having multiple stand-alone firewalls.
security services	See cryptography .
serial transmission	A method of data transmission in which the bits of a data character are transmitted sequentially over a single channel.
SGCP	Simple Gateway Control Protocol. Controls VoIP gateways by an external call control element (called a call-agent).
SGSN	Serving GPRS Support Node. The SGSN ensures mobility management, session management and packet relaying functions.
SHA-1	Secure Hash Algorithm 1. SHA-1 [NIST94c] is a revision to SHA that was published in 1994. SHA is closely modeled after MD4 and produces a 160-bit digest. Because SHA produces a 160-bit digest, it is more resistant to brute-force attacks than 128-bit hashes (such as MD5), but it is slower. Secure Hash Algorithm 1 is a joint creation of the National Institute of Standards and Technology and the National Security Agency. This algorithm, like other hash algorithms, is used to generate a hash value, also known as a message digest, that acts like a CRC used in lower-layer protocols to ensure that message contents are not changed during transmission. SHA-1 is generally considered more secure than MD5 .
SIP	Session Initiation Protocol. Enables call handling sessions, particularly two-party audio conferences, or “calls.” SIP works with SDP for call signaling. SDP specifies the ports for the media stream. Using SIP, the security appliance can support any SIP VoIP gateways and VoIP proxy servers.
site-to-site VPN	A site-to-site VPN is established between two IPSec peers that connect remote networks into a single VPN . In this type of VPN , neither IPSec peer is the destination or source of user traffic. Instead, each IPSec peer provides encryption and authentication services for hosts on the LAN s connected to each IPSec peer. The hosts on each LAN send and receive data through the secure tunnel established by the pair of IPSec peers.
SKEME	A key exchange protocol that defines how to derive authenticated keying material, with rapid key refreshment.
SMR	Stub Multicast Routing. SMR allows the security appliance to function as a “stub router.” A stub router is a device that acts as an IGMP proxy agent. IGMP is used to dynamically register specific hosts in a multicast group on a particular LAN with a multicast router. Multicast routers route multicast data transmissions to hosts that are registered to receive specific multimedia or other broadcasts. A stub router forwards IGMP messages between hosts and MC routers .
SMTP	Simple Mail Transfer Protocol. SMTP is an Internet protocol that supports email services.
SNMP	Simple Network Management Protocol. A standard method for managing network devices using data structures called Management Information Bases.

split tunneling	Allows a remote VPN client simultaneous encrypted access to a private network and clear unencrypted access to the Internet . If you do not enable split tunneling, all traffic between the VPN client and the security appliance is sent through an IPSec tunnel. All traffic originating from the VPN client is sent to the outside interface through a tunnel, and client access to the Internet from its remote site is denied.
spoofing	A type of attack designed to foil network security mechanisms such as filters and access lists. A spoofing attack sends a packet that claims to be from an address from which it was not actually sent.
SQL*Net	Structured Query Language Protocol. An Oracle protocol used to communicate between client and server processes.
SSH	Secure Shell. An application running on top of a reliable transport layer, such as TCP/IP, that provides strong authentication and encryption capabilities.
SSL	Secure Sockets Layer. A protocol that resides between the application layer and TCP/IP to provide transparent encryption of data traffic.
standby unit	See secondary unit .
stateful inspection	Network protocols maintain certain data, called state information, at each end of a network connection between two hosts. State information is necessary to implement the features of a protocol, such as guaranteed packet delivery, data sequencing, flow control, and transaction or session IDs. Some of the protocol state information is sent in each packet while each protocol is being used. For example, a browser connected to a web server uses HTTP and supporting TCP/IP protocols. Each protocol layer maintains state information in the packets it sends and receives. The security appliance and some other firewalls inspect the state information in each packet to verify that it is current and valid for every protocol it contains. This is called stateful inspection and is designed to create a powerful barrier to certain types of computer security threats.
Static PAT	Static Port Address Translation. Static PAT is a static address that also maps a local port to a global port. See also Dynamic PAT , NAT .
subnetmask	See mask .

T

TACACS+	Terminal Access Controller Access Control System Plus. A client-server protocol that supports AAA services, including command authorization. See also AAA , RADIUS .
TAPI	Telephony Application Programming Interface. A programming interface in Microsoft Windows that supports telephony functions.
TCP	Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission.

TCP Intercept	With the TCP intercept feature, once the optional embryonic connection limit is reached, and until the embryonic connection count falls below this threshold, every SYN bound for the effected server is intercepted. For each SYN, the security appliance responds on behalf of the server with an empty SYN/ACK segment. The security appliance retains pertinent state information, drops the packet, and waits for the client acknowledgment. If the ACK is received, then a copy of the client SYN segment is sent to the server and the TCP three-way handshake is performed between the security appliance and the server. If this three-way handshake completes, may the connection resume as normal. If the client does not respond during any part of the connection phase, then the security appliance retransmits the necessary segment using exponential back-offs.
TDP	Tag Distribution Protocol. TDP is used by tag switching devices to distribute, request, and release tag binding information for multiple network layer protocols in a tag switching network. TDP does not replace routing protocols. Instead, it uses information learned from routing protocols to create tag bindings. TDP is also used to open, monitor, and close TDP sessions and to indicate errors that occur during those sessions. TDP operates over a connection-oriented transport layer protocol with guaranteed sequential delivery (such as TCP). The use of TDP does not preclude the use of other mechanisms to distribute tag binding information, such as piggybacking information on other protocols.
Telnet	A terminal emulation protocol for TCP/IP networks such as the Internet . Telnet is a common way to control web servers remotely; however, its security vulnerabilities have led to its replacement by SSH .
TFTP	Trivial File Transfer Protocol. TFTP is a simple protocol used to transfer files. It runs on UDP and is explained in depth in RFC 1350.
TID	Tunnel Identifier.
TLS	Transport Layer Security. A future IETF protocol to replace SSL .
traffic policing	The traffic policing feature ensures that no traffic exceeds the maximum rate (bits per second) that you configure, thus ensuring that no one traffic flow can take over the entire resource.
transform set	See IPSec transform set .
translate, translation	See xlate .
transparent firewall mode	A mode in which the security appliance is not a router hop. You can use transparent firewall mode to simplify your network configuration or to make the security appliance invisible to attackers. You can also use transparent firewall mode to allow traffic through that would otherwise be blocked in routed firewall mode . See also routed firewall mode .
transport mode	An IPSec encryption mode that encrypts only the data portion (payload) of each packet, but leaves the header untouched. Transport mode is less secure than tunnel mode.
TSP	TAPI Service Provider. See also TAPI .
tunnel mode	An IPSec encryption mode that encrypts both the header and data portion (payload) of each packet. Tunnel mode is more secure than transport mode.

tunnel	A method of transporting data in one protocol by encapsulating it in another protocol. Tunneling is used for reasons of incompatibility, implementation simplification, or security. For example, a tunnel lets a remote VPN client have encrypted access to a private network.
Turbo ACL	Increases ACL lookup speeds by compiling them into a set of lookup tables. Packet headers are used to access the tables in a small, fixed number of lookups, independent of the existing number of ACL entries.

U

UDP	User Datagram Protocol. A connectionless transport layer protocol in the IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, which requires other protocols to handle error processing and retransmission. UDP is defined in RFC 768.
UMTS	Universal Mobile Telecommunication System. An extension of GPRS networks that moves toward an all-IP network by delivering broadband information, including commerce and entertainment services, to mobile users via fixed, wireless, and satellite networks
Unicast RPF	Unicast Reverse Path Forwarding. Unicast RPF guards against spoofing by ensuring that packets have a source IP address that matches the correct source interface according to the routing table.
URL	Uniform Resource Locator. A standardized addressing scheme for accessing hypertext documents and other services using a browser. For example, http://www.cisco.com .
user EXEC mode	User EXEC mode lets you to see the security appliance settings. The user EXEC mode prompt appears as follows when you first access the security appliance. See also command-specific configuration mode , global configuration mode , and privileged EXEC mode .
UTC	Coordinated Universal Time. The time zone at zero degrees longitude, previously called Greenwich Mean Time (GMT) and Zulu time. UTC replaced GMT in 1967 as the world time standard. UTC is based on an atomic time scale rather than an astronomical time scale.
UTRAN	Universal Terrestrial Radio Access Network. Networking protocol used for implementing wireless networks in UMTS. GTP allows multi-protocol packets to be tunneled through a UMTS/GPRS backbone between a GGSN , an SGSN and the UTRAN .
UUIE	User-User Information Element. An element of an H.225 packet that identifies the users implicated in the message.

V

VLAN	Virtual LAN . A group of devices on one or more LANs that are configured (using management software) so that they can communicate as if they were attached to the same physical network cable, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.
VoIP	Voice over IP. VoIP carries normal voice traffic, such as telephone calls and faxes, over an IP-based network. DSP segments the voice signal into frames, which then are coupled in groups of two and stored in voice packets. These voice packets are transported using IP in compliance with ITU-T specification H.323 .

VPN Virtual Private Network. A network connection between two peers over the public network that is made private by strict authentication of users and the encryption of all data traffic. You can establish VPNs between clients, such as PCs, or a [headend](#), such as the security appliance.

virtual firewall See [security context](#).

VSA Vendor-specific attribute. An attribute in a [RADIUS](#) packet that is defined by a vendor rather than by [RADIUS](#) RFCs. The [RADIUS](#) protocol uses IANA-assigned vendor numbers to help identify VSAs. This lets different vendors have VSAs of the same number. The combination of a vendor number and a VSA number makes a VSA unique. For example, the cisco-av-pair VSA is attribute 1 in the set of VSAs related to vendor number 9. Each vendor can define up to 256 VSAs. A [RADIUS](#) packet contains any VSAs attribute 26, named Vendor-specific. VSAs are sometimes referred to as subattributes.

W

WAN wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers.

WCCP Web Cache Communication Protocol. Transparently redirects selected types of traffic to a group of web cache engines to optimize resource usage and lower response times.

Websense A content filtering solution that manages employee access to the [Internet](#). Websense uses a policy engine and a [URL](#) database to control user access to websites.

WEP Wired Equivalent Privacy. A security protocol for wireless [LANs](#), defined in the IEEE 802.11b standard.

WINS Windows Internet Naming Service. A Windows system that determines the IP address associated with a particular network device, also known as “name resolution.” WINS uses a distributed database that is automatically updated with the [NetBIOS](#) names of network devices currently available and the IP address assigned to each one. WINS provides a distributed database for registering and querying dynamic [NetBIOS](#) names to IP address mapping in a routed network environment. It is the best choice for [NetBIOS](#) name resolution in such a routed network because it is designed to solve the problems that occur with name resolution in complex networks.

X

X.509 A widely used standard for defining digital certificates. X.509 is actually an ITU recommendation, which means that it has not yet been officially defined or approved for standardized usage.

xauth See [IKE Extended Authentication](#).

xlate An xlate, also referred to as a translation entry, represents the mapping of one IP address to another, or the mapping of one IP address/port pair to another.



INDEX

Symbols

/bits subnet masks [D-3](#)
?
 command string [C-4](#)
 help [C-4](#)

Numerics

4GE SSM
 connector types [5-1](#)
 fiber [5-3](#)
 SFP [5-3](#)
 support [A-9](#)
802.1Q tagging [4-11](#)
802.1Q trunk [5-3](#)

A

AAA
 about [13-1](#)
 accounting [19-13](#)
 addressing, configuring [31-2](#)
 authentication
 CLI access [40-5](#)
 network access [19-1](#)
 privileged EXEC mode [40-6](#)
 authorization
 command [40-7](#)
 downloadable access lists [19-8](#)
 network access [19-6](#)
 local database support [13-9](#)
 performance [19-1](#)

server
 adding [13-12](#)
 types [13-2](#)
 support summary [13-3](#)
 web clients [19-5](#)
abbreviating commands [C-3](#)
Access Control Server [33-1, 33-2, 33-5](#)
access hours, username attribute [30-72](#)
accessing the security appliance using SSL [37-3](#)
accessing the security appliance using TKS1 [37-3](#)
access list filter, username attribute [30-74](#)
access lists
 about [16-1](#)
 ACE logging, configuring [16-20](#)
 comments [16-18](#)
 deny flows, managing [16-22](#)
 downloadable [19-9](#)
 EtherType, adding [16-8](#)
 exemptions from posture validation [33-4](#)
 extended
 about [16-5](#)
 adding [16-6](#)
 group policy WebVPN filter [30-64](#)
 implicit deny [16-3](#)
 inbound [18-1](#)
 interface, applying [18-2](#)
 IP address guidelines [16-3](#)
 IPsec [27-20](#)
 logging [16-20](#)
 NAT guidelines [16-3](#)
 Network Admission Control, default [33-3](#)
 object groups [16-18](#)
 outbound [18-1](#)

- remarks [16-18](#)
- scheduling activation [16-18](#)
- standard, adding [16-11](#)
- types [16-2](#)
- WebVPN username connections [30-81](#)
- access ports [4-9](#)
- ACEs
 - See* access lists
- Active/Active failover
 - about [14-10](#)
 - actions [14-14](#)
 - command replication [14-12](#)
 - configuration synchronization [14-11](#)
 - configuring
 - asymmetric routing support [14-35](#)
 - cable-based failover [14-27](#)
 - failover criteria [14-34](#)
 - failover group preemption [14-33](#)
 - HTTP replication [14-34](#)
 - interface monitoring [14-34](#)
 - LAN-based failover [14-29](#)
 - prerequisites [14-27](#)
 - virtual MAC addresses [14-35](#)
 - device initialization [14-11](#)
 - duplicate MAC addresses, avoiding [14-11, 14-35](#)
 - primary status [14-11](#)
 - secondary status [14-11](#)
 - triggers [14-13](#)
- Active/Standby failover
 - about [14-6](#)
 - actions [14-9](#)
 - command replication [14-8](#)
 - configuration synchronization [14-7](#)
 - configuring
 - cable-based [14-20](#)
 - failover criteria [14-26](#)
 - HTTP replication [14-25](#)
 - interface monitoring [14-25](#)
 - interface poll times [14-39](#)
 - LAN-based [14-21](#)
 - prerequisites [14-20](#)
 - unit poll times [14-39](#)
 - virtual MAC addresses [14-26](#)
 - device initialization [14-7](#)
 - primary unit [14-7](#)
 - secondary unit [14-7](#)
 - triggers [14-9](#)
- Active Directory, settings for password management [30-24](#)
- Active Directory procedures [E-19 to E-22](#)
- Adaptive Security Algorithm [1-4](#)
- admin context
 - about [3-2](#)
 - changing [6-13](#)
- administrative distance [9-3](#)
- Advanced Encryption Standard (AES) [27-3](#)
- AIP SSM
 - about [22-1](#)
 - checking status [22-13](#)
 - configuration [22-2](#)
 - initial setup [22-4](#)
 - loading an image [22-14](#)
 - sending traffic to [22-2](#)
 - support [A-9](#)
- alternate address, ICMP message [D-15](#)
- Application Access Panel, WebVPN [37-33](#)
- application access using WebVPN
 - and e-mail proxy [37-50](#)
 - and hosts file errors [37-18](#)
 - and Web Access [37-50](#)
 - configuring client applications [37-49](#)
 - enabling cookies on browser [37-49](#)
 - group policy WebVPN attribute [30-65](#)
 - privileges [37-49](#)
 - quitting properly [37-19](#)
 - re-enabling [37-20](#)
 - setting up on client [37-49](#)
 - username WebVPN attribute [30-82](#)

- using e-mail [37-50](#)
 - with IMAP client [37-50](#)
- application inspection
 - about [25-2](#)
 - applying [25-5](#)
 - configuring [25-5](#)
 - inspection class map [21-11](#)
 - inspection policy map [21-8](#)
 - security level requirements [7-1](#)
 - special actions [21-7](#)
- Application Profile Customization Framework [37-30](#)
- ARP inspection
 - about [26-1](#)
 - enabling [26-2](#)
 - static entry [26-2](#)
- ARP spoofing [26-2](#)
- ARP test, failover [14-17](#)
- ASA (Adaptive Security Algorithm) [1-4](#)
- ASA 5505
 - Base license [4-2](#)
 - client
 - authentication [34-11](#)
 - configuration restrictions, table [34-2](#)
 - device pass-through [34-8](#)
 - group policy attributes pushed to [34-9](#)
 - mode [34-3](#)
 - remote management [34-8](#)
 - split tunneling [34-7](#)
 - TCP [34-4](#)
 - trustpoint [34-7](#)
 - tunnel group [34-6](#)
 - tunneling [34-5](#)
 - Xauth [34-4](#)
 - interfaces, about [4-1](#)
 - MAC addresses [4-4](#)
 - maximum VLANs [4-2](#)
 - native VLAN support [4-11](#)
 - non-forwarding interface [4-6](#)
 - power over Ethernet [4-4](#)
 - protected switch ports [4-9](#)
 - Security Plus license [4-2](#)
 - server (headend) [34-1](#)
 - SPAN [4-4](#)
 - Spanning Tree Protocol, unsupported [4-9](#)
 - VLAN interface configuration [4-5](#)
- ASDM software
 - allowing access [40-3](#)
 - configuring ASDM and WebVPN on the same interface [40-4](#)
 - installing [41-3](#)
- ASR [14-35](#)
- asymmetric routing support [14-35](#)
- attributes
 - LDAP [E-5](#)
 - policy [E-2](#)
 - RADIUS [E-25](#)
 - username [30-72](#)
- attribute-value pairs
 - TACACS+ [E-32](#)
- attribute-value pairs (AVP) [30-31](#)
- authentication
 - about [13-1](#)
 - ASA 5505 as Easy VPN client [34-11](#)
 - CLI access [40-5](#)
 - FTP [19-3](#)
 - HTTP [19-2](#)
 - network access [19-1](#)
 - privileged EXEC mode [40-6](#)
 - restrictions, WebVPN [37-5](#)
 - Telnet [19-2](#)
 - web clients [19-5](#)
 - WebVPN users with digital certificates [37-15](#)
- authorization
 - about [13-2](#)
 - command [40-7](#)
 - downloadable access lists [19-8](#)
 - network access [19-6](#)
- Auto-MDI/MDIX [5-1](#)

auto-signon

group policy WebVPN attribute [30-64](#)

username WebVPN attribute [30-83](#)

Auto-Update, configuring [41-10](#)

B

backup device, load balancing [29-5](#)

backup server attributes, group policy [30-48](#)

Baltimore Technologies, CA server support [39-5](#)

banner message, group policy [30-41](#)

bits subnet masks [D-3](#)

Black Ice firewall [30-57](#)

BPDUs

ACL, EtherType [16-10](#)

BPDUs, EtherType access list [16-10](#)

bridge

entry timeout [26-4](#)

table, *See* MAC address table

broadcast Ping test [14-18](#)

bypass authentication [34-8](#)

C

CA

certificate validation, not done in WebVPN [37-2](#)

CRs and [39-2](#)

public key cryptography [39-1](#)

revoked certificates [39-2](#)

server support [39-5](#)

supported servers [39-5](#)

caching [37-28](#)

capturing packets [43-11](#)

cascading access lists [27-15](#)

certificate

authentication, e-mail proxy [37-27](#)

enrollment protocol [39-7](#)

group matching

configuring [27-9](#)

rule and policy, creating [27-10](#)

Certificate Revocation Lists

See CRLs

certification authority

See CA

changing between contexts [6-11](#)

Cisco-AV-Pair LDAP attributes [E-14](#)

Cisco Integrated Firewall [30-56](#)

Cisco IP Phones

DHCP [10-4](#)

Cisco IP Phones, application inspection [25-71](#)

Cisco LDAP attributes [E-5](#)

Cisco Security Agent [30-56](#)

Cisco Trust Agent [33-5](#)

Class A, B, and C addresses [D-1](#)

class-default class map [21-4](#)

classes, logging

filtering messages by [42-16](#)

message class variables [42-16](#)

types [42-16](#)

classes, MPF

See class map

classes, resource

See resource management

class map

inspection [21-11](#)

Layer 3/4

management traffic [21-7](#)

match commands [21-5](#)

through traffic [21-5](#)

regular expression [21-14](#)

CLI

abbreviating commands [C-3](#)

adding comments [C-6](#)

command line editing [C-3](#)

command output paging [C-5](#)

displaying [C-5](#)

help [C-4](#)

- paging [C-5](#)
 - syntax formatting [C-3](#)
- client
 - VPN 3002 hardware, forcing client update [29-3](#)
 - Windows, client update notification [29-3](#)
- client access rules, group policy [30-58](#)
- client firewall, group policy [30-55](#)
- clientless authentication [33-5](#)
- client mode [34-3](#)
- client update, performing [29-3](#)
- cluster
 - IP address, load balancing [29-6](#)
 - load balancing configurations [29-7](#)
 - mixed scenarios [29-8](#)
 - virtual [29-5](#)
- command authorization
 - about [40-7](#)
 - configuring [40-7](#)
- command prompts [C-2](#)
- comments
 - access lists [16-18](#)
 - configuration [C-6](#)
- configuration
 - clearing [2-9](#)
 - comments [C-6](#)
 - factory default
 - commands [2-1](#)
 - restoring [2-2](#)
 - saving [2-6](#)
 - text file [2-9](#)
 - URL for a context [6-9](#)
 - viewing [2-8](#)
- configuration mode
 - accessing [2-5](#)
 - prompt [C-2](#)
- connection blocking [23-11](#)
- connection limits
 - configuring [23-6](#)
 - per context [6-6](#)
- connect time, maximum, username attribute [30-74](#)
- console port logging [42-8](#)
- content transformation, WebVPN [37-28](#)
- contexts
 - See* security contexts
- conversion error, ICMP message [D-16](#)
- cookies, enabling for WebVPN [37-5](#)
- CRACK protocol [27-28](#)
- crash dump [43-11](#)
- crypto map
 - access lists [27-20](#)
 - applying to interfaces [27-20, 36-7](#)
 - clearing configurations [27-27](#)
 - creating an entry to use the dynamic crypto map [32-7](#)
 - definition [27-12](#)
 - dynamic [27-24](#)
 - dynamic, creating [32-6](#)
 - entries [27-12](#)
 - examples [27-21](#)
 - policy [27-13](#)
- crypto show commands [27-26](#)
- CSC SSM
 - about [22-5](#)
 - checking status [22-13](#)
 - failover [22-7](#)
 - getting started [22-7](#)
 - loading an image [22-14](#)
 - sending traffic to [22-11](#)
 - support [A-9](#)
 - what to scan [22-9](#)
- CSD support [A-11](#)
- custom firewall [30-56](#)
- customization, WebVPN
 - group policy WebVPN attribute [30-61](#)
 - login windows for WebVPN users [30-23](#)
 - username WebVPN attribute [30-21, 30-80](#)
- cut-through proxy [19-1](#)

D

data flow

- routed firewall [15-3](#)
- transparent firewall [15-13](#)

DDNS [10-6](#)

debugging IPSec [28-7](#)

debug messages [43-11](#)

default

- class [6-3](#)
- DefaultL2Lgroup [30-1](#)
- DefaultRAGroup [30-1](#)
- domain name, group policy [30-43](#)
- group policy [30-1, 30-31](#)
- LAN-to-LAN tunnel group [30-13](#)
- remote access tunnel group, configuring [30-5](#)
- routes, defining equal cost routes [9-4](#)
- tunnel group [27-11, 30-2](#)

default configuration

- commands [2-1](#)
- restoring [2-2](#)

default policy [21-3](#)

default routes

- about [9-4](#)
- configuring [9-4](#)

deny flows, logging [16-22](#)

deny in a crypto map [27-15](#)

deny-message

- group policy WebVPN attribute [30-62](#)
- username WebVPN attribute [30-81](#)

DES, IKE policy keywords (table) [27-3](#)

device ID, including in messages [42-19](#)

device pass-through, ASA 5505 as Easy VPN client [34-8](#)

DfltGrpPolicy [30-32](#)

DHCP

- addressing, configuring [31-3](#)
- Cisco IP Phones [10-4](#)
- options [10-3](#)
- relay [10-5](#)

server [10-1, 10-2](#)

transparent firewall [16-6](#)

DHCP Intercept, configuring [30-44](#)

Diffie-Hellman

Group 5 [27-4](#)

groups supported [27-4](#)

DiffServ preservation [24-5](#)

digital certificates

- authenticating WebVPN users [37-15](#)
- SSL [37-5](#)
- WebVPN authentication restrictions [37-5](#)

directory hierarchy search [E-4](#)

disabling content rewrite [37-29](#)

disabling messages, specific message IDs [42-20](#)

DMZ, definition [1-1](#)

DNS

configuring for WebVPN [37-16](#)

dynamic [10-6](#)

inspection

- about [25-13](#)
- managing [25-13](#)
- rewrite, about [25-14](#)
- rewrite, configuring [25-15](#)

NAT effect on [17-14](#)

server, configuring [30-35](#)

domain attributes, group policy [30-43](#)

domain name [8-2](#)

dotted decimal subnet masks [D-3](#)

downloadable access lists

- configuring [19-9](#)
- converting netmask expressions [19-12](#)

DSCP preservation [24-5](#)

dual IP stack, configuring [12-4](#)

dual-ISP support [9-5](#)

duplex, configuring [5-1](#)

dynamic crypto map [27-24](#)

creating [32-6](#)

See also crypto map

Dynamic DNS [10-6](#)

dynamic NAT

See NAT

E

Easy VPN

client

authentication [34-11](#)

configuration restrictions, table [34-2](#)

enabling and disabling [34-1](#)

group policy attributes pushed to [34-9](#)

mode [34-3](#)

remote management [34-8](#)

trustpoint [34-7](#)

tunnels [34-8](#)

Xauth [34-4](#)

server (headend) [34-1](#)

Easy VPN client

ASA 5505

device pass-through [34-8](#)

split tunneling [34-7](#)

TCP [34-4](#)

tunnel group [34-6](#)

tunneling [34-5](#)

echo reply, ICMP message [D-15](#)

ECMP [9-3](#)

editing command lines [C-3](#)

EIGRP [16-6](#)

e-mail

closing the Outlook connection [37-27](#)

configuring for WebVPN [37-26](#)

proxies, WebVPN [37-26](#)

proxy, certificate authentication [37-27](#)

WebVPN, configuring [37-26](#)

EMBLEM format, using in logs [42-20](#)

enable command [2-5](#)

end-user interface, WebVPN, defining [37-32](#)

Enterprises [10-4](#)

Entrust, CA server support [39-5](#)

established command, security level requirements [7-2](#)

Ethernet

Auto-MDI/MDIX [5-1](#)

duplex [5-1](#)

speed [5-1](#)

EtherType

assigned numbers [16-10](#)

See also access lists

external group policy, configuring [30-34](#)

F

facility, syslog [42-8](#)

factory default configuration

commands [2-1](#)

restoring [2-2](#)

failover

about [14-1](#)

Active/Active, configuring [14-27](#)

Active/Active, *See* Active/Active failover

Active/Standby, configuring [14-19](#)

Active/Standby, *See* Active/Standby failover

configuration file

terminal messages, Active/Active [14-12](#)

terminal messages, Active/Standby [14-7](#)

configuring [14-19](#)

contexts [14-7](#)

controlling [14-49](#)

debug messages [14-51](#)

disabling [14-50](#)

displaying commands [14-48](#)

encrypting failover communication [14-39](#)

Ethernet failover cable [14-4](#)

examples

Active/Active LAN-based failover [B-22, B-28](#)

Active/Standby cable-based failover [B-20, B-26](#)

Active/Standby LAN-based failover [B-21, B-27](#)

failover link [14-3](#)

forcing [14-49](#)

- health monitoring [14-16](#)
 - interface health [14-17](#)
 - interface monitoring [14-17](#)
 - interface tests [14-17](#)
 - licenses [14-2](#)
 - link communications [14-3](#)
 - MAC addresses
 - about [14-7](#)
 - automatically assigning [6-11](#)
 - monitoring, configuration [14-49](#)
 - monitoring, health [14-16](#)
 - network tests [14-17](#)
 - primary unit [14-7](#)
 - restoring a failed group [14-50](#)
 - restoring a failed unit [14-50](#)
 - secondary unit [14-7](#)
 - serial cable [14-4](#)
 - SNMP syslog traps [14-51](#)
 - software versions [14-2](#)
 - Stateful Failover, *See* Stateful Failover
 - state link [14-5](#)
 - subsecond [14-39](#)
 - system log messages [14-51](#)
 - system requirements [14-2](#)
 - testing [14-49](#)
 - type selection [14-15](#)
 - understanding [14-1](#)
 - unit health [14-17](#)
 - verifying the configuration [14-40](#)
- fast path [1-4](#)
- fiber interfaces [5-3](#)
- filter (access list)
 - group policy WebVPN attribute [30-64](#)
 - username WebVPN attribute [30-81](#)
- filtering
 - about [20-1](#)
 - ActiveX [20-2](#)
 - FTP [20-9](#)
 - Java applets [20-3](#)
 - security level requirements [7-2](#)
 - servers supported [20-4](#)
 - show command output [C-4](#)
 - URLs [20-4](#)
- firewall
 - Black Ice [30-57](#)
 - Cisco Integrated [30-56](#)
 - Cisco Security Agent [30-56](#)
 - custom [30-56](#)
 - Network Ice [30-57](#)
 - none [30-56](#)
 - Sygate personal [30-57](#)
 - Zone Labs [30-56](#)
- firewall mode
 - about [15-1](#)
 - configuring [2-5](#)
- firewall policy, group policy [30-55](#)
- FO (failover) license [14-3](#)
- FO_AA license [14-3](#)
- format of messages [42-23](#)
- fragmentation policy, IPsec [27-8](#)
- fragment size [23-11](#)
- FTP inspection
 - about [25-27](#)
 - configuring [25-26](#)
- functions, WebVPN
 - username WebVPN attribute [30-77](#)
 - WebVPN group policy attribute [30-60](#)

G

- general attributes, tunnel group [30-2](#)
- general parameters, tunnel group [30-2](#)
- general tunnel-group connection parameters [30-2](#)
- generating RSA keys [39-6](#)
- global addresses
 - recommendations [17-14](#)
 - specifying [17-24](#)
- global e-mail proxy attributes [37-26](#)

- global IPsec SA lifetimes, changing [27-22](#)
- group-lock, username attribute [30-75](#)
- group policy
 - address pools [30-54](#)
 - attributes [30-35](#)
 - backup server attributes [30-48](#)
 - client access rules [30-58](#)
 - configuring [30-34](#)
 - default domain name for tunneled packets [30-43](#)
 - definition [30-1, 30-31](#)
 - domain attributes [30-43](#)
 - Easy VPN client, attributes pushed to ASA 5505 [34-9](#)
 - external, configuring [30-34](#)
 - firewall policy [30-55](#)
 - hardware client user idle timeout [30-46](#)
 - internal, configuring [30-35](#)
 - IP phone bypass [30-46](#)
 - IPSec over UDP attributes [30-41](#)
 - LEAP Bypass [30-47](#)
 - network extension mode [30-47](#)
 - security attributes [30-39](#)
 - split tunneling attributes [30-42](#)
 - split-tunneling domains [30-44](#)
 - user authentication [30-45](#)
 - VPN attributes [30-36](#)
 - VPN hardware client attributes [30-45](#)
 - webvpn attributes [30-59](#)
 - WINS and DNS servers [30-35](#)
- group policy, default [30-31](#)
- group policy, secure unit authentication [30-45](#)
- group policy WebVPN attributes
 - application access [30-65](#)
 - auto-signon [30-64](#)
 - customization [30-61](#)
 - deny-message [30-62](#)
 - filter [30-64](#)
 - home page [30-63](#)
 - html-content filter [30-62](#)
 - keep-alive-ignore [30-66](#)

- port forward [30-65](#)
- port-forward-name [30-66](#)
- sso-server [30-67](#)
- svc [30-68](#)
- url-list [30-65](#)

- GTP inspection
 - about [25-32](#)
 - configuring [25-32](#)

H

- H.225 timeouts [25-42](#)
- H.245 troubleshooting [25-43](#)
- H.323 inspection
 - about [25-38](#)
 - configuring [25-38](#)
 - limitations [25-39](#)
 - troubleshooting [25-44](#)
- hairpinning [27-20](#)
- hardware client, group policy attributes [30-45](#)
- help, command line [C-4](#)
- HMAC hashing method [27-3](#)
- hold-period [33-8](#)
- homepage
 - group policy WebVPN attribute [30-63](#)
 - username WebVPN attribute [30-80](#)
- hostname
 - configuring [8-2](#)
 - in banners [8-2](#)
 - multiple context mode [8-2](#)
- hosts, subnet masks for [D-3](#)
- hosts file
 - errors [37-18](#)
 - reconfiguring [37-20](#)
 - WebVPN [37-19](#)
- HSRP [15-9](#)
- html-content-filter
 - group policy WebVPN attribute [30-62](#)
 - username WebVPN attribute [30-79](#)

HTTP(S)

- authentication [40-5](#)
- filtering [20-4](#)

HTTP/HTTPS Web VPN proxy, setting [37-5](#)

HTTP compression, WebVPN, enabling [30-67, 30-84](#)

HTTP inspection

- about [25-44](#)
- configuring [25-44](#)

HTTP redirection for login, Easy VPN client on the ASA 5505 [34-12](#)

HTTPS for WebVPN sessions [37-3](#)

hub-and-spoke VPN scenario [27-20](#)

ICMP

- testing connectivity [43-1](#)
- type numbers [D-15](#)

idle timeout

- hardware client user, group policy [30-46](#)
- username attribute [30-73](#)

ID method for ISAKMP peers, determining [27-6](#)

IKE

- benefits [27-2](#)
- creating policies [27-4](#)
- keepalive setting, tunnel group [30-3](#)
- pre-shared key, Easy VPN client on the ASA 5505 [34-6](#)
- See also* ISAKMP

ILS inspection [25-53](#)

IM [25-65](#)

inbound access lists [18-1](#)

Individual user authentication [34-12](#)

information reply, ICMP message [D-15](#)

information request, ICMP message [D-15](#)

inheritance

- tunnel group [30-1](#)
- username attribute [30-72](#)

inside, definition [1-1](#)

inspection_default class-map [21-4](#)

inspection engines

See application inspection

Instant Messaging inspection [25-65](#)

intercept DHCP, configuring [30-44](#)

interfaces

ASA 5505

- about [4-1](#)
- enabled status [4-9](#)
- IP address [4-7](#)
- MAC addresses [4-4](#)
- maximum VLANs [4-2](#)
- non-forwarding [4-6](#)
- protected switch ports [4-9](#)
- switch port configuration [4-9](#)
- trunk ports [4-11](#)
- VLAN interface configuration [4-5](#)

configuring for remote access [32-2](#)

configuring IPv6 on [12-3](#)

duplex [5-1](#)

enabled status [5-1](#)

enabling [5-2](#)

failover monitoring [14-17](#)

fiber [5-3](#)

global addresses [17-24](#)

IDs [5-2](#)

IP address [7-4](#)

MAC addresses

- automatically assigning [6-11](#)
- manually assigning to interfaces [7-4](#)

mapped name [6-8](#)

naming, physical and subinterface [7-3](#)

naming, VLAN [4-6](#)

SFP [5-3](#)

speed [5-1](#)

subinterfaces [5-3](#)

viewing monitored interface status [14-48](#)

internal group policy, configuring [30-35](#)

Internet Security Association and Key Management Protocol

See ISAKMP

intrusion prevention configuration [22-2](#)

IP addresses

ASA 5505 [4-7](#)

classes [D-1](#)

configuring an assignment method for remote access clients [31-1](#)

configuring for VPNs [31-1](#)

configuring local IP address pools [31-2](#)

interface [7-4](#)

management, transparent firewall [8-5](#)

private [D-2](#)

subnet mask [D-4](#)

IP phone [34-8](#)

IP phone bypass, group policy [30-46](#)

IPS configuration [22-2](#)

IPSec

anti-replay window [24-12](#)

enabling debug [28-7](#)

modes [28-2](#)

over UDP, group policy, configuring attributes [30-41](#)

remote-access tunnel group [30-6](#)

setting maximum active VPN sessions [29-3](#)

IPsec

access list [27-20](#)

basic configuration with static crypto maps [27-22](#)

Cisco VPN Client [27-2](#)

configuring [27-1, 27-11](#)

crypto map entries [27-12](#)

fragmentation policy [27-8](#)

over NAT-T, enabling [27-7](#)

over TCP, enabling [27-8](#)

SA lifetimes, changing [27-22](#)

tunnel [27-11](#)

viewing configuration [27-26](#)

IPsec parameters, tunnel group [30-3](#)

ipsec-ra, creating an IPsec remote-access tunnel [30-6](#)

IP spoofing, preventing [23-10](#)

IPv6

access lists [12-6](#)

commands [12-1](#)

configuring alongside IPv4 [12-4](#)

default route [12-5](#)

dual IP stack [12-4](#)

duplicate address detection [12-4](#)

enabling [12-3](#)

neighbor discovery [12-7](#)

router advertisement messages [12-9](#)

static neighbor [12-11](#)

static routes [12-5](#)

verifying [12-11](#)

IPv6 addresses

anycast [D-9](#)

command support for [12-1](#)

format [D-5](#)

multicast [D-8](#)

prefixes [D-10](#)

required [D-10](#)

types of [D-6](#)

unicast [D-6](#)

ISAKMP

about [27-2](#)

configuring [27-1, 27-2](#)

determining an ID method for peers [27-6](#)

disabling in aggressive mode [27-6](#)

enabling on the outside interface [27-6, 32-3](#)

keepalive setting, tunnel group [30-3](#)

policies, configuring [27-5](#)

See also IKE

J

Java applets, filtering [20-2](#)

Java object signing [37-29](#)

java-trustpoint [37-29](#)

K

keep-alive-ignore

- group policy WebVPN attribute [30-66](#)
- username WebVPN attribute [30-83](#)

Kerberos

- configuring [13-12](#)
- support [13-5](#)

L

L2TP description [28-1](#)

LAN-to-LAN tunnel group, configuring [30-13](#)

latency

- about [24-1](#)
- configuring [24-2, 24-3](#)
- reducing [24-7](#)

Layer 2 firewall

- See* transparent firewall

Layer 2 forwarding table

- See* MAC address table

Layer 2 Tunneling Protocol [28-1](#)

Layer 3/4

- matching multiple policy maps [21-18](#)

LDAP

- AAA support [13-6](#)
- application inspection [25-53](#)
- attribute mapping [13-8](#)
- Cisco attributes [E-5](#)
- Cisco-AV-pair [E-14](#)
- configuring [13-12](#)
- configuring a AAA server [E-2 to E-18](#)
- directory about [E-3](#)
- directory search [E-4](#)
- example configuration procedures [E-19 to E-22](#)
- hierarchy example [E-3](#)
- permissions policy [E-2](#)
- SASL [13-6](#)
- schema example [E-15](#)

schema loading [E-18](#)

schema planning [E-3 to E-5](#)

server configuration about [E-3](#)

server type [13-7](#)

user authentication [13-6](#)

user authorization [13-7](#)

user permissions [E-18](#)

LEAP Bypass, group policy [30-47](#)

licenses

- FO [14-3](#)
- FO_AA [14-3](#)
- managing [41-1](#)
- per model [A-1](#)
- UR [14-3](#)

link up/down test [14-17](#)

LLQ

- See* low-latency queue

load balancing

- cluster configurations [29-7](#)
- concepts [29-5](#)
- eligible clients [29-7](#)
- eligible platforms [29-7](#)
- implementing [29-6](#)
- mixed cluster scenarios [29-8](#)
- platforms [29-7](#)
- prerequisites [29-6](#)

local user database

- adding a user [13-11](#)
- configuring [13-10](#)
- logging in [40-6](#)
- support [13-9](#)

lockout recovery [40-15](#)

log buffer

- save to internal Flash [42-13](#)
- send to FTP server [42-14](#)

logging

- access lists [16-20](#)
- classes
 - filtering messages by [42-15](#)

- types [42-16](#)
- device-id, including in system log messages [42-19](#)
- e-mail
 - configuring as output destination [42-9](#)
 - destination address [42-10](#)
 - source address [42-9](#)
- EMBLEM format [42-20](#)
- facility option [42-8](#)
- filtering
 - by message class [42-16](#)
 - by message list [42-17](#)
 - by severity level [42-5](#)
- logging queue, configuring [42-19](#)
- output destinations
 - ASDM [42-10](#)
 - console port [42-8](#)
 - email address [42-9](#)
 - internal buffer [42-5](#)
 - SNMP [42-4](#)
 - syslog server [42-7](#)
 - Telnet or SSH session [42-5](#)
- queue
 - changing the size of [42-19](#)
 - configuring [42-19](#)
 - viewing queue statistics [42-19](#)
- severity level
 - changing [42-21](#)
- severity level, changing [42-21](#)
- timestamp, including [42-19](#)
- login
 - banner, configuring [40-16](#)
 - console [2-5](#)
 - enable [2-5](#)
 - FTP [19-3](#)
 - global configuration mode [2-5](#)
 - local user [40-6](#)
 - password [8-1](#)
 - simultaneous, username attribute [30-73](#)
 - SSH [40-3](#)

- Telnet [8-1](#)
- windows, customizing for WebVPN users [30-23](#)
- low-latency queue
 - applying [24-2, 24-3](#)

M

- MAC addresses
 - ASA 5505 [4-4](#)
 - ASA 5505 device pass-through [34-8](#)
 - automatically assigning [6-11](#)
 - failover [14-7](#)
 - manually assigning to interfaces [7-4](#)
 - security context classification [3-3](#)
- MAC address table
 - about [15-13](#)
 - built-in-switch [26-3](#)
 - entry timeout [26-4](#)
 - MAC learning, disabling [26-4](#)
 - resource management [6-6](#)
 - static entry [26-3](#)
- MAC learning, disabling [26-4](#)
- management IP address, transparent firewall [8-5](#)
- man-in-the-middle attack [26-2](#)
- MAPI, configuring [37-27](#)
- mapped interface name [6-8](#)
- mask
 - reply, ICMP message [D-16](#)
 - request, ICMP message [D-15](#)
- match commands
 - inspection class map [21-9](#)
 - Layer 3/4 class map [21-5](#)
- matching, certificate group [27-9](#)
- maximum active IPsec VPN sessions, setting [29-3](#)
- maximum connect time, username attribute [30-74](#)
- maximum object size to ignore username WebVPN attribute [30-83](#)
- maximum sessions, IPsec [29-11](#)
- MD5, IKE policy keywords (table) [27-3](#)

- message list
 - filtering by [42-17](#)
- message-of-the-day banner [40-16](#)
- messages, logging
 - classes
 - about [42-15](#)
 - list of [42-16](#)
 - component descriptions [42-23](#)
 - filtering by message list [42-17](#)
 - format of [42-23](#)
 - message list, creating [42-17](#)
 - severity levels [42-23](#)
- metacharacters, regular expression [21-12](#)
- MGCP inspection
 - about [25-56](#)
 - configuring [25-56](#)
- MIBs [42-1](#)
- Microsoft Active Directory, settings for password management [30-24](#)
- Microsoft Internet Explorer client parameters, configuring [30-49](#)
- Microsoft Windows 2000 CA, supported [39-5](#)
- mixed cluster scenarios, load balancing [29-8](#)
- mobile redirection, ICMP message [D-16](#)
- mode
 - context [3-10](#)
 - firewall [2-5](#)
- Modular Policy Framework
 - See* MPF
- monitoring
 - failover [14-16](#)
 - OSPF [9-19](#)
 - resource management [6-16](#)
 - SNMP [42-1](#)
- monitoring switch traffic, ASA 5505 [4-4](#)
- More prompt [C-5](#)
- MPF
 - about [21-1](#)
 - default policy [21-3](#)
 - examples [21-21](#)
 - feature directionality [21-17](#)
 - features [21-1](#)
 - flows [21-18](#)
 - matching multiple policy maps [21-18](#)
 - service policy, applying [21-21](#)
 - See also* class map
 - See also* policy map
- MPLS
 - LDP [16-10](#)
 - router-id [16-10](#)
 - TDP [16-10](#)
- MSIE client parameters, configuring [30-49](#)
- MTU size, Easy VPN client, ASA 5505 [34-5](#)
- multicast traffic [15-10](#)
- multiple context mode
 - See* security contexts

N

- NAC
 - See* Network Admission Control
- naming an interface
 - ASA 5505 [4-6](#)
 - other models [7-3](#)
- NAT
 - about [17-1, 17-2](#)
 - bypassing NAT
 - about [17-9](#)
 - configuration [17-30](#)
 - DNS [17-14](#)
 - dynamic NAT
 - about [17-5](#)
 - configuring [17-23](#)
 - implementation [17-17](#)
 - examples [17-33](#)
 - exemption from NAT
 - about [17-9](#)
 - configuration [17-32](#)

- identity NAT
 - about [17-9](#)
 - configuration [17-30](#)
- NAT ID [17-17](#)
- order of statements [17-14](#)
- overlapping addresses [17-34](#)
- PAT
 - about [17-7](#)
 - configuring [17-23](#)
 - implementation [17-17](#)
- policy NAT
 - about [17-9](#)
- port redirection [17-35](#)
- RPC not supported with [25-77](#)
- same security level [17-13](#)
- security level requirements [7-2](#)
- static identify, configuring [17-30](#)
- static NAT
 - about [17-7](#)
 - configuring [17-26](#)
- static PAT
 - about [17-8](#)
 - configuring [17-27](#)
- transparent firewall [15-12](#)
- types [17-5](#)
- native VLAN support [4-11](#)
- NAT-T
 - enabling IPsec over NAT-T [27-7](#)
 - using [27-7](#)
- Netscape CMS, CA server support [39-5](#)
- Network Activity test [14-17](#)
- Network Address Translation
 - See* NAT
- Network Admission Control
 - Access Control Server [33-2](#)
 - ACL, default [33-3](#)
 - clientless authentication [33-5](#)
 - configuring [30-51](#)
 - enabling and disabling [33-2](#)

- exemptions [33-4](#)
- port [33-7](#)
- retransmission retries [33-7](#)
- retransmission retry timer [33-7](#)
- revalidation timer [33-9](#)
- session reinitialization timer [33-8](#)
- uses, requirements, and limitations [33-1](#)
- network extension mode [34-3](#)
- network extension mode, group policy [30-47](#)
- Network Ice firewall [30-57](#)
- networks, overlapping [17-34](#)
- Nokia VPN Client [27-28](#)
- NTLM support [13-5](#)
- NT server
 - configuring [13-12](#)
 - support [13-5](#)

O

- object groups
 - nesting [16-15](#)
 - removing [16-17](#)
- open ports [D-14](#)
- operating systems, posture validation exemptions [33-4](#)
- OSPF
 - about [9-9](#)
 - area authentication [9-14](#)
 - area MD5 authentication [9-14](#)
 - area parameters [9-13](#)
 - authentication key [9-11](#)
 - cost [9-12](#)
 - dead interval [9-12](#)
 - default route [9-17](#)
 - displaying update packet pacing [9-19](#)
 - enabling [9-10](#)
 - hello interval [9-12](#)
 - interface parameters [9-11](#)
 - link-state advertisement [9-9](#)
 - logging neighbor states [9-18](#)

- MD5 authentication [9-12](#)
- monitoring [9-19](#)
- NSSA [9-14](#)
- packet pacing [9-19](#)
- processes [9-9](#)
- redistributing routes [9-10](#)
- route calculation timers [9-17](#)
- route map [9-7](#)
- route summarization [9-15](#)
- stub area [9-14](#)
- summary route cost [9-14](#)
- outbound access lists [18-1](#)
- Outlook connection, closing [37-27](#)
- Outlook Exchange proxy, configuring [37-27](#)
- Outlook Web Access (OWA) and WebVPN [37-50](#)
- output destinations [42-5](#)
 - e-mail address [42-5](#), [42-9](#)
 - SNMP management station [42-5](#)
 - specifying [42-9](#)
 - syslog server [42-5](#), [42-7](#)
 - Telnet or SSH session [42-5](#)
 - viewing logs [42-7](#)
- outside, definition [1-1](#)
- oversubscribing resources [6-2](#)

P

- packet
 - capture [43-11](#)
 - classifier [3-3](#)
- packet flow
 - routed firewall [15-3](#)
 - transparent firewall [15-13](#)
- paging screen displays [C-5](#)
- parameter problem, ICMP message [D-15](#)
- password
 - resetting on SSM hardware module [43-10](#)
- password management, Active Directory settings [30-24](#)
- passwords
 - changing [8-1](#)
 - clientless authentication [33-6](#)
 - recovery [43-7](#)
 - security appliance [8-1](#)
 - username, setting [30-71](#)
 - WebVPN [37-43](#)
- password-storage, username attribute [30-76](#)
- PAT
 - Easy VPN client mode [34-3](#)
 - See also* NAT
 - static [17-27](#)
- PDA support for WebVPN [37-25](#)
- peers
 - alerting before disconnecting [27-9](#)
 - ISAKMP, determining ID method [27-6](#)
- performance, optimizing for WebVPN [37-28](#)
- permit in a crypto map [27-15](#)
- ping
 - See* ICMP
- PKI protocol [39-7](#)
- PoE [4-4](#)
- policing
 - flow within a tunnel [24-9](#)
- policy, QoS [24-1](#)
- policy map
 - inspection [21-8](#)
 - Layer 3/4
 - about [21-15](#)
 - adding [21-19](#)
 - default policy [21-18](#)
 - feature directionality [21-17](#)
 - flows [21-18](#)
- policy NAT
 - about [17-9](#)
 - dynamic, configuring [17-24](#)
 - static, configuring [17-26](#)
 - static PAT, configuring [17-28](#)
- pools, address
 - DHCP [10-2](#)

- global NAT [17-24](#)
- port-forward
 - group policy WebVPN attribute [30-65](#)
 - username WebVPN attribute [30-82](#)
- port forwarding
 - automatic applet download [37-18](#)
 - configuring client applications [37-49](#)
- port-forward-name
 - group policy WebVPN attribute [30-66](#)
 - username WebVPN attribute [30-83](#)
- ports
 - open on device [D-14](#)
 - redirection, NAT [17-35](#)
 - TCP and UDP [D-11](#)
- posture validation
 - exemptions [33-4](#)
 - port [33-7](#)
 - revalidation timer [33-9](#)
 - uses, requirements, and limitations [33-1](#)
- power over Ethernet [4-4](#)
- PPPoE, configuring [35-1 to 35-5](#)
- pre-shared key, Easy VPN client on the ASA 5505 [34-6](#)
- primary unit, failover [14-7](#)
- printers [34-8](#)
- private networks [D-2](#)
- privileged EXEC mode, accessing [2-5](#)
- privileged mode
 - accessing [2-5](#)
 - prompt [C-2](#)
- privilege level, username, setting [30-71](#)
- prompts
 - command [C-2](#)
 - more [C-5](#)
- protocol numbers and literal values [D-11](#)
- proxy
 - See* e-mail proxy
- proxy bypass [37-29](#)
- proxy servers
 - SIP and [25-65](#)

- public key cryptography [39-1](#)

Q

QoS

- about [24-1, 24-3](#)
- DiffServ preservation [24-5](#)
- DSCP preservation [24-5](#)
- feature interaction [24-4](#)
- policies [24-1](#)
- priority queueing
 - IPSec anti-replay window [24-12](#)
- statistics [24-13](#)
- token bucket [24-2](#)
- traffic shaping
 - overview [24-4](#)
- viewing statistics [24-13](#)

Quality of Service

See QoS

question mark

- command string [C-4](#)
- help [C-4](#)

queue, logging

- changing the size of [42-19](#)
- viewing statistics [42-19](#)

queue, QoS

- latency, reducing [24-7](#)
- limit [24-2, 24-3](#)

R

RADIUS

- attribute policy [E-2](#)
- attributes [E-25](#)
- Cisco AV pair [E-14](#)
- configuring a AAA server [E-24](#)
- configuring a server [13-12](#)
- downloadable access lists [19-9](#)

- network access authentication [19-3](#)
- network access authorization [19-8](#)
- permissions policy [E-2](#)
- support [13-3](#)
- RAS, H.323 troubleshooting [25-44](#)
- rate limiting [24-3](#)
- RealPlayer [25-63](#)
- reboot, waiting until active sessions end [27-9](#)
- redirect, ICMP message [D-15](#)
- redundancy, in site-to-site VPNs, using crypto maps [27-26](#)
- Registration Authority description [39-2](#)
- regular expression [21-12](#)
- reloading
 - context [6-14](#)
 - security appliance [43-6](#)
- remarks [16-18](#)
- remote access
 - configuration summary [32-1](#)
 - IPSec tunnel group, configuring [30-6](#)
 - restricting [30-75](#)
 - tunnel group, configuring default [30-5](#)
 - user, adding [32-4](#)
 - VPN, configuring [32-1](#)
- remote management, ASA 5505 [34-8](#)
- resetting the SSM hardware module password [43-10](#)
- resource management
 - about [6-2](#)
 - assigning a context [6-10](#)
 - class [6-4](#)
 - configuring [6-1](#)
 - default class [6-3](#)
 - monitoring [6-16](#)
 - oversubscribing [6-2](#)
 - resource types [6-6](#)
 - unlimited [6-2](#)
- resource usage [6-19](#)
- retransmission retries, Network Admission Control [33-7](#)
- revalidation timer, Network Admission Control [33-9](#)
- revoked certificates [39-2](#)
- rewrite, disabling [37-29](#)
- RIP
 - about [9-20](#)
 - enabling [9-20](#)
- routed mode
 - about [15-1](#)
 - setting [2-5](#)
- route maps
 - defining [9-7](#)
 - uses [9-7](#)
- router
 - advertisement, ICMP message [D-15](#)
 - solicitation, ICMP message [D-15](#)
- routes
 - about default [9-4](#)
 - about static [9-2](#)
 - configuring default routes [9-4](#)
 - configuring IPv6 default [12-5](#)
 - configuring IPv6 static [12-5](#)
 - configuring static routes [9-3](#)
- routing
 - OSPF [9-20](#)
 - other protocols [16-6](#)
- RS-232 cable
 - See* failover [14-4](#)
- RSA
 - KEON, CA server support [39-5](#)
 - keys, generating [39-6, 40-2](#)
 - signatures, IKE authentication method [39-2](#)
- RTSP inspection
 - about [25-63](#)
 - configuring [25-63](#)
- running configuration
 - copying [41-8](#)
 - saving [2-6](#)

S

same security level communication

enabling [7-6](#)

NAT [17-13](#)

SAs, lifetimes [27-22](#)

SCCP (Skinny) inspection

about [25-71](#)

configuration [25-71](#)

configuring [25-71](#)

SDI

configuring [13-12](#)

support [13-4](#)

secondary device, virtual cluster [29-5](#)

secondary unit, failover [14-7](#)

secure unit authentication [34-11](#)

secure unit authentication, group policy [30-45](#)

security, WebVPN [37-2, 37-6](#)

Security Agent, Cisco [30-56](#)

security appliance

CLI [C-1](#)

connecting to [2-4](#)

managing licenses [41-1](#)

managing the configuration [2-6](#)

reloading [43-6](#)

upgrading software [41-3](#)

viewing files in Flash memory [41-2](#)

security association

clearing [27-27](#)

See also SAs

security attributes, group policy [30-39](#)

security contexts

about [3-1](#)

adding [6-7](#)

admin context

about [3-2](#)

changing [6-13](#)

assigning to a resource class [6-10](#)

cascading [3-8](#)

changing between [6-11](#)

classifier [3-3](#)

configuration

URL, changing [6-13](#)

URL, setting [6-9](#)

logging in [3-9](#)

MAC addresses

automatically assigning [6-11](#)

classifying using [3-3](#)

managing [6-1, 6-12](#)

mapped interface name [6-8](#)

monitoring [6-15](#)

multiple mode, enabling [3-10](#)

nesting or cascading [3-9](#)

prompt [C-2](#)

reloading [6-14](#)

removing [6-12](#)

resource management [6-2](#)

resource usage [6-19](#)

saving all configurations [2-7](#)

unsupported features [3-2](#)

VLAN allocation [6-7](#)

security level

about [7-1](#)

interface [7-3](#)

interface, ASA 5505 [4-6](#)

serial cable

See failover

server group [33-2](#)

service policy

applying [21-21](#)

default [21-21](#)

global [21-21](#)

interface [21-21](#)

session management path [1-4](#)

session reinitialization timer, Network Admission Control [33-8](#)

severity levels, of system log messages

changing [42-5](#)

- filtering by [42-5](#)
 - list of [42-23](#)
- severity levels, of system messages
 - definition [42-23](#)
- SHA, IKE policy keywords (table) [27-3](#)
- show command, filtering output [C-4](#)
- simultaneous logins, username attribute [30-73](#)
- single mode
 - backing up configuration [3-10](#)
 - configuration [3-10](#)
 - enabling [3-10](#)
 - restoring [3-11](#)
- single sign-on
 - See* SSO
- single-signon
 - group policy WebVPN attribute [30-67](#)
 - username WebVPN attribute [30-85](#)
- SIP inspection
 - about [25-65](#)
 - configuring [25-65](#)
 - instant messaging [25-65](#)
 - timeouts [25-70](#)
 - troubleshooting [25-70](#)
- site-to-site VPNs, redundancy [27-26](#)
- SMTP inspection [25-74](#)
- SNMP
 - about [42-1](#)
 - management station [42-5](#)
 - MIBs [42-1](#)
 - traps [42-2](#)
- source quench, ICMP message [D-15](#)
- SPAN [4-4](#)
- Spanning Tree Protocol, unsupported [4-9](#)
- speed, configuring [5-1](#)
- split tunneling
 - ASA 5505 as Easy VPN client [34-7](#)
 - group policy [30-42](#)
 - group policy, domains [30-44](#)
- SSH
 - authentication [40-5](#)
 - concurrent connections [40-2](#)
 - login [40-3](#)
 - password [8-1](#)
 - RSA key [40-2](#)
 - username [40-3](#)
- SSL
 - certificate [37-5](#)
 - used to access the security appliance [37-3](#)
- SSL/TLS encryption protocols
 - configuring [37-4](#)
 - WebVPN [37-4](#)
- SSL VPN Client
 - benefits [38-1](#)
 - compression [38-6](#)
 - DPD [38-5](#)
 - enabling [38-3](#)
 - address assignment [38-3](#)
 - groups and users [38-4](#)
 - permanent installation [38-4](#)
 - tunnel group [38-3](#)
 - group policy WebVPN attribute [30-68](#)
 - installing [38-2](#)
 - images [38-2](#)
 - order [38-2](#)
 - keepalive messages [38-6](#)
 - logging out sessions [38-8](#)
 - username WebVPN attribute [30-85](#)
 - viewing sessions [38-7](#)
- SSM
 - checking status [22-13](#)
 - configuration
 - AIP SSM [22-2](#)
 - CSC SSM [22-7](#)
 - loading an image [22-14](#)
 - See also* AIP SSM
 - See also* CSC SSM
- sso-server
 - group policy WebVPN attribute [30-67](#)

- username WebVPN attribute [30-85](#)
- SSO with WebVPN [37-6 to 37-15](#)
 - configuring HTTP Basic and NTLM authentication [37-6](#)
 - configuring HTTP form protocol [37-9](#)
 - configuring SiteMinder [37-7](#)
- startup configuration
 - copying [41-8](#)
 - saving [2-6](#)
- Stateful Failover
 - about [14-16](#)
 - state information [14-16](#)
 - state link [14-5](#)
 - statistics [14-43, 14-47](#)
- stateful inspection [1-4](#)
- state information [14-16](#)
- state link [14-5](#)
- static ARP entry [26-2](#)
- static bridge entry [26-3](#)
- static NAT
 - See* NAT
- static PAT
 - See* PAT
- static routes
 - about [9-2](#)
 - configuring [9-3](#)
 - tracking [9-5](#)
- statistics, QoS [24-13](#)
- stealth firewall
 - See* transparent firewall
- subcommand mode prompt [C-2](#)
- subinterfaces, adding [5-3](#)
- subnet masks
 - /bits [D-3](#)
 - about [D-2](#)
 - address range [D-4](#)
 - determining [D-3](#)
 - dotted decimal [D-3](#)
 - number of hosts [D-3](#)
- Sun Microsystems Java Runtime Environment and WebVPN [37-49](#)
- Sun RPC inspection
 - about [25-77](#)
 - configuring [25-77](#)
- SVC
 - See* SSL VPN Client
- svc
 - group policy WebVPN attribute [30-68](#)
 - username WebVPN attribute [30-85](#)
- switch MAC address table [26-3](#)
- switch ports
 - access ports [4-9](#)
 - default configuration [4-4](#)
 - protected [4-9](#)
 - SPAN [4-4](#)
 - trunk ports [4-11](#)
- Sygate Personal Firewall [30-57](#)
- SYN attacks, monitoring [6-20](#)
- SYN cookies [6-20](#)
- syntax formatting [C-3](#)
- syslog server
 - as output destination [42-7](#)
 - designating [42-7](#)
 - designating more than one [42-8](#)
- EMBLEM format
 - configuring [42-20](#)
 - enabling [42-7](#)
- system configuration [3-2](#)
- system log messages
 - classes [42-16](#)
 - classes of [42-15](#)
 - configuring in groups
 - by message list [42-17](#)
 - by severity level [42-5](#)
 - creating lists of [42-15](#)
 - device ID, including [42-19](#)
 - disabling logging of [42-5](#)
 - filtering by message class [42-15](#)

- managing in groups
 - by message class [42-16](#)
 - creating a message list [42-15](#)
- output destinations [42-5](#)
 - email address [42-9](#)
 - SNMP [42-4](#)
 - syslog message server [42-5](#)
 - Telnet or SSH session [42-5](#)
- severity levels
 - about [42-23](#)
 - changing the severity level of a message [42-5](#)
- timestamp, including [42-19](#)

T

TACACS+

- command authorization, configuring [40-11](#)
- configuring a server [13-12](#)
- network access authorization [19-7](#)
- support [13-4](#)

tail drop [24-3](#)

TCP

- ASA 5505 as Easy VPN client [34-4](#)
- connection limits per context [6-6](#)
- ports and literal values [D-11](#)
- sequence number randomization
 - disabling in NAT configuration [17-24](#)
 - disabling using Modular Policy Framework [23-8](#)

TCP Intercept

- enabling using Modular Policy Framework [23-8](#)
- enabling using NAT [17-24](#)
- monitoring [6-20](#)

TCP normalization [23-1](#)

Telnet

- allowing management access [40-1](#)
- authentication [40-5](#)
- concurrent connections [40-1](#)
- password [8-1](#)

testing configuration [43-1](#)

- time exceeded, ICMP message [D-15](#)
- time ranges, access lists [16-18](#)
- timestamp, including in system log messages [42-19](#)
- timestamp reply, ICMP message [D-15](#)
- timestamp request, ICMP message [D-15](#)
- TLS1, used to access the security appliance [37-3](#)
- token bucket [24-2](#)
- toolbar, floating, WebVPN [37-34](#)
- traffic flow
 - routed firewall [15-3](#)
 - transparent firewall [15-13](#)
- traffic shaping
 - overview [24-4](#)
- Transform [27-12](#)
- transform set
 - creating [32-4](#)
 - definition [27-12](#)
- transmit queue ring limit [24-2, 24-3](#)
- transparent firewall
 - about [15-8](#)
 - ARP inspection
 - about [26-1](#)
 - enabling [26-2](#)
 - static entry [26-2](#)
 - data flow [15-13](#)
 - DHCP packets, allowing [16-6](#)
 - guidelines [15-10](#)
 - HSRP [15-9](#)
 - MAC address timeout [26-4](#)
 - MAC learning, disabling [26-4](#)
 - Management 0/0 IP address [7-4](#)
 - management IP address [8-5](#)
 - multicast traffic [15-10](#)
 - NAT [15-12](#)
 - packet handling [16-6](#)
 - static bridge entry [26-3](#)
 - unsupported features [15-11](#)
 - VRRP [15-9](#)
- traps, SNMP [42-2](#)

- troubleshooting
 - H.323 [25-43](#)
 - H.323 RAS [25-44](#)
 - SIP [25-70](#)
 - trunk, 802.1Q [5-3](#)
 - trunk ports [4-11](#)
 - trustpoint [39-3](#)
 - trustpoint, ASA 5505 client [34-7](#)
 - tunnel
 - ASA 5505 as Easy VPN client [34-5](#)
 - IPsec [27-11](#)
 - security appliance as a tunnel endpoint [27-1](#)
 - tunnel group
 - ASA 5505 as Easy VPN client [34-6](#)
 - configuring [30-5](#)
 - creating [30-6](#)
 - default [27-11, 30-1, 30-2](#)
 - default, remote access, configuring [30-5](#)
 - default LAN-to-LAN, configuring [30-13](#)
 - definition [30-1, 30-2](#)
 - general parameters [30-2](#)
 - inheritance [30-1](#)
 - IPSec parameters [30-3](#)
 - LAN-to-LAN, configuring [30-13](#)
 - name and type [30-6](#)
 - remote access, configuring [32-5](#)
 - remote-access, configuring [30-6](#)
 - tunnel-group
 - general attributes [30-2](#)
 - webvpn attributes [30-4](#)
 - tunnel-group ISAKMP/IKE keepalive settings [30-3](#)
 - tunneling, about [27-1](#)
 - tunnel mode [28-2](#)
 - tx-ring-limit [24-2, 24-3](#)
 - connection state information [1-5](#)
 - ports and literal values [D-11](#)
 - unreachable, ICMP message [D-15](#)
 - UR (unrestricted) license [14-2](#)
 - url-list
 - group policy WebVPN attribute [30-65](#)
 - username WebVPN attribute [30-82](#)
 - URLs
 - context configuration, changing [6-13](#)
 - context configuration, setting [6-9](#)
 - filtering, about [20-4](#)
 - filtering, configuration [20-6](#)
 - user, VPN
 - definition [30-1](#)
 - remote access, adding [32-4](#)
 - user access, restricting remote [30-75](#)
 - user authentication, group policy [30-45](#)
 - user EXEC mode
 - accessing [2-5](#)
 - prompt [C-2](#)
 - username
 - adding [13-10](#)
 - clientless authentication [33-6](#)
 - encrypted [13-11](#)
 - management tunnels [34-8](#)
 - password [13-11](#)
 - WebVPN [37-43](#)
 - Xauth for Easy VPN client [34-4](#)
 - username attributes
 - access hours [30-72](#)
 - configuring [30-70, 30-72](#)
 - group-lock [30-75](#)
 - inheritance [30-72](#)
 - password, setting [30-71](#)
 - password-storage [30-76](#)
 - privilege level, setting [30-71](#)
 - simultaneous logins [30-73](#)
 - vpn-filter [30-74](#)
 - vpn-framed-ip-address [30-74](#)
-
- ## U
- UDP
 - connection limits per context [6-6](#)

- vpn-idle timeout [30-73](#)
- vpn-session-timeout [30-74](#)
- vpn-tunnel-protocol [30-75](#)
- username configuration, viewing [30-71](#)
- username WebVPN attributes
 - auto-signon [30-83](#)
 - customization [30-21](#), [30-80](#)
 - deny message [30-81](#)
 - filter (access list) [30-81](#)
 - functions [30-77](#)
 - homepage [30-80](#)
 - html-content-filter [30-79](#)
 - keep-alive ignore [30-83](#)
 - port-forward [30-82](#)
 - port-forward-name [30-83](#)
 - sso-server [30-85](#)
 - svc [30-85](#)
 - url-list [30-82](#)
- username WebVPN mode [30-76](#)
- U-turn [27-20](#)

V

- VeriSign, configuring CAs example [39-5](#)
- viewing logs [42-7](#)
- viewing QoS statistics [24-13](#)
- viewing RMS [41-13](#)
- virtual cluster [29-5](#)
 - IP address [29-6](#)
 - master [29-5](#)
- virtual firewalls
 - See* security contexts
- virtual HTTP [19-3](#)
- VLANs [5-4](#)
 - 802.1Q trunk [5-3](#)
 - allocating to a context [6-7](#)
 - ASA 5505
 - configuring [4-5](#)
 - MAC addresses [4-4](#)

- maximum [4-2](#)
- mapped interface name [6-8](#)
- subinterfaces [5-3](#)

VoIP

- proxy servers [25-65](#)
- troubleshooting [25-43](#)

VPN

- address pool, configuring [32-4](#)
- address pool, configuring (group-policy) [30-54](#)
- address range, subnets [D-4](#)
- parameters, general, setting [29-1](#)
- setting maximum number of IPSec sessions [29-3](#)

- VPN attributes, group policy [30-36](#)

- VPN Client, IPsec attributes [27-2](#)

- vpn-filter username attribute [30-74](#)

- vpn-framed-ip-address username attribute [30-74](#)

- VPN hardware client, group policy attributes [30-45](#)

- vpn-idle-timeout username attribute [30-73](#)

- vpn load balancing

- See* load balancing [29-5](#)

- vpn-session-timeout username attribute [30-74](#)

- vpn-tunnel-protocol username attribute [30-75](#)

- VRRP [15-9](#)

W

- WCCP [10-9](#)

- web browsing with WebVPN [37-47](#)

- web caching [10-9](#)

- web clients, secure authentication [19-5](#)

- web e-Mail (Outlook Web Access), Outlook Web Access [37-27](#)

WebVPN

- assigning users to group policies [37-16](#)
- authenticating with digital certificates [37-15](#)
- CA certificate validation not done [37-2](#)
- client application requirements [37-44](#)
- client requirements [37-44](#)
 - for file management [37-48](#)

- for network browsing [37-48](#)
- for port forwarding [37-49](#)
- for using applications [37-49](#)
- for web browsing [37-47](#)
- start-up [37-46](#)
- configuring
 - DNS globally [37-16](#)
 - e-mail [37-26](#)
- configuring for specific users [30-76](#)
- configuring WebVPN and ASDM on the same interface [37-3](#)
- cookies [37-5](#)
- defining the end-user interface [37-32](#)
- definition [37-1](#)
- digital certificate authentication restrictions [37-5](#)
- e-mail [37-26](#)
- e-mail proxies [37-26](#)
- enable cookies for [37-49](#)
- end user set-up [37-32](#)
- establishing a session [37-3](#)
- floating toolbar [37-34](#)
- group policy attributes, configuring [37-17](#)
- hosts file [37-19](#)
- hosts files, reconfiguring [37-20](#)
- HTTP/HTTPS proxy, setting [37-5](#)
- Java object signing [37-29](#)
- PDA support [37-25](#)
- printing and [37-46](#)
- remote system configuration and end-user requirements [37-46](#)
- security precautions [37-2](#), [37-6](#)
- security tips [37-44](#)
- setting HTTP/HTTPS proxy [37-4](#)
- SSL/TLS encryption protocols [37-4](#)
- supported applications [37-44](#)
- supported browsers [37-46](#)
- supported types of Internet connections [37-46](#)
- troubleshooting [37-18](#)
- unsupported features [37-2](#)

- URL [37-46](#)
- use of HTTPS [37-3](#)
- username and password required [37-46](#)
- usernames and passwords [37-43](#)
- use suggestions [37-32](#), [37-44](#)
- WebVPN, Application Access Panel [37-33](#)
- webvpn attributes
 - group policy [30-59](#)
 - tunnel-group [30-4](#)
- WebVPN group policy attributes functions [30-60](#)
- WebVPN tunnel-group connection parameters [30-4](#)
- welcome message, group policy [30-41](#)
- WINS server, configuring [30-35](#)

X

- Xauth, Easy VPN client [34-4](#)

Z

- Zone Labs firewalls [30-56](#)
- Zone Labs Integrity Server [13-16](#)

