

ASDM Book 2: Cisco ASA Series Firewall ASDM Configuration Guide, 7.16

First Published: 2021-01-29

Last Modified: 2026-07-17

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

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.

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

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

About This Guide xix

Document Objectives xix

Related Documentation xix

Document Conventions xix

Communications, Services, and Additional Information xxi

CHAPTER 1

Introduction to ASA-Firewall Services 1

How to Implement Firewall Services 1

Basic Access Control 2

Application Filtering 2

URL Filtering 2

Threat Protection 3

Firewall Services for Virtual Environments 4

Network Address Translation 4

Application Inspection 5

Use Case: Expose a Server to the Public 5

PART I

Access Control 9

CHAPTER 2

Access Rules 11

Controlling Network Access 11

General Information About Rules 12

Interface Access Rules and Global Access Rules 12

Inbound and Outbound Rules 12

Rule Order 13

Implicit Permits 13

Implicit Deny 14
NAT and Access Rules 14
Same Security Level Interfaces and Access Rules 14
Extended Access Rules 15
Extended Access Rules for Returning Traffic 15
Allowing Broadcast and Multicast Traffic 15
Management Access Rules 15
EtherType Rules 16
Supported EtherTypes and Other Traffic 16
EtherType Rules for Returning Traffic 16
Allowing MPLS 16
Licensing for Access Rules 17
Guidelines for Access Control 17
Configure Access Control 19
Configure Access Rules 19
Access Rule Properties 20
Configure Advanced Options for Access Rules 22
Configure Management Access Rules 23
Configure EtherType Rules 24
Configure ICMP Access Rules 26
Monitoring Access Rules 27
Evaluating Syslog Messages for Access Rules 27
History for Access Rules 28

CHAPTER 3 **Objects for Access Control** 31
Guidelines for Objects 31
Configure Objects 32
Configure Network Objects and Groups 32
Configure a Network Object 32
Configure a Network Object Group 33
Configure Service Objects and Service Groups 33
Configure a Service Object 33
Configure a Service Group 34
Configure Local User Groups 35

Configure Security Group Object Groups 36
Configure Time Ranges 37
Monitoring Objects 38
History for Objects 38

CHAPTER 4 **Access Control Lists 41**
About ACLs 41
ACL Types 41
The ACL Manager 43
ACL Names 43
Access Control Entry Order 43
Permit/Deny vs. Match/Do Not Match 44
Access Control Implicit Deny 44
IP Addresses Used for Extended ACLs When You Use NAT 44
Time-Based ACEs 45
Licensing for Access Control Lists 45
Guidelines for ACLs 46
Configure ACLs 47
Configure Extended ACLs 47
Extended ACE Properties 48
Service Specifications in Extended ACEs 50
Configure Standard ACLs 51
Configure Webtype ACLs 52
Webtype ACE Properties 53
Examples for Webtype ACLs 54
Monitoring ACLs 54
History for ACLs 55

CHAPTER 5 **Identity Firewall 57**
About the Identity Firewall 57
Architecture for Identity Firewall Deployments 58
Features of the Identity Firewall 59
Deployment Scenarios 60
Guidelines for the Identity Firewall 63

Prerequisites for the Identity Firewall 65
Configure the Identity Firewall 66
Configure the Active Directory Domain 66
Configure Active Directory Server Groups 67
Configure Active Directory Agents 67
Configure Active Directory Agent Groups 68
Configure Identity Options 68
Configure Identity-Based Security Policy 71
Monitoring the Identity Firewall 72
History for the Identity Firewall 72

CHAPTER 6 ASA and Cisco TrustSec 73
About Cisco TrustSec 73
About SGT and SXP Support in Cisco TrustSec 74
Roles in the Cisco TrustSec Feature 74
Security Group Policy Enforcement 75
How the ASA Enforces Security Group-Based Policies 76
Effects of Changes to Security Groups on the ISE 77
Speaker and Listener Roles on the ASA 78
Register the ASA with the ISE 79
Create a Security Group on the ISE 80
Generate the PAC File 80
Guidelines for Cisco TrustSec 81
Configure the ASA to Integrate with Cisco Trustsec 83
Configure the AAA Server for Cisco TrustSec Integration 84
Import a PAC File 85
Configure the Security Exchange Protocol 86
Add an SXP Connection Peer 87
Refresh Environment Data 88
Configure the Security Policy 89
Configure Layer 2 Security Group Tagging Imposition 89
Usage Scenarios 90
Configure a Security Group Tag on an Interface 91
Configure IP-SGT Bindings Manually 92

AnyConnect Client VPN Support for Cisco TrustSec 92
Add an SGT to Remote Access VPN Group Policies and Local Users 93
Monitoring Cisco TrustSec 94
History for Cisco TrustSec 95

CHAPTER 7 **ASA FirePOWER Module 97**
About the ASA FirePOWER Module 97
How the ASA FirePOWER Module Works with the ASA 97
ASA FirePOWER Inline Mode 98
ASA FirePOWER Inline Tap Monitor-Only Mode 99
ASA FirePOWER Passive Monitor-Only Traffic Forwarding Mode 99
ASA FirePOWER Management 100
Compatibility with ASA Features 100
What to Do if the ASA FirePOWER Module Cannot Filter URLs 100
Licensing Requirements for the ASA FirePOWER Module 101
Guidelines for ASA FirePOWER 101
Defaults for ASA FirePOWER 103
Perform Initial ASA FirePOWER Setup 103
Deploy the ASA FirePOWER Module in Your Network 103
Routed Mode 103
Transparent Mode 105
Register the ASA FirePOWER Module with a Management Center 106
Access the ASA FirePOWER CLI 106
Configure ASA FirePOWER Basic Settings 107
Configure the ASA FirePOWER Module for ASDM Management 108
Configure the ASA FirePOWER Module 111
Configure the Security Policy on the ASA FirePOWER Module 111
Redirect Traffic to the ASA FirePOWER Module 111
Configure Inline or Inline Tap Monitor-Only Modes 111
Configure Passive Traffic Forwarding 112
Enable Captive Portal for Active Authentication 113
Managing the ASA FirePOWER Module 114
Install or Reimage the Module 115
Install or Reimage the Software Module 115

Reset the Password 118
Reload or Reset the Module 118
Shut Down the Module 119
Uninstall a Software Module Image 119
Session to the Software Module From the ASA 120
Upgrade the System Software 120
Monitoring the ASA FirePOWER Module 121
Showing Module Status 121
Showing Module Statistics 121
Analyzing Operational Behavior (ASDM Management) 121
Monitoring Module Connections 122
History for the ASA FirePOWER Module 123

CHAPTER 8 **Cisco Umbrella** 125
About Cisco Umbrella Connector 125
Cisco Umbrella Enterprise Security Policy 125
Cisco Umbrella Registration 126
Licensing Requirements for Cisco Umbrella Connector 126
Guidelines and Limitations for Cisco Umbrella 126
Configure Cisco Umbrella Connector 128
Install the CA Certificate from the Cisco Umbrella Registration Server 129
Configure the Umbrella Connector Global Settings 130
Enable Umbrella in the DNS Inspection Policy Map 131
Verify the Umbrella Registration 132
Monitoring the Umbrella Connector 133
Monitoring the Umbrella Service Policy Statistics 133
Monitoring Umbrella Syslog Messages 135
History for Cisco Umbrella Connector 136

PART II **Firewall Services for Virtual Environments** 137

CHAPTER 9 **Attribute-Based Access Control** 139
Guidelines for Attribute-Based Network Objects 139
Configure Attribute-Based Access Control 140

Configure Attributes for vCenter Virtual Machines 140
Configure a VM Attribute Agent 142
Configure Attribute-Based Network Objects 143
Configure Access Rules Using Attribute-Based Network Objects 144
Monitoring Attribute-Based Network Objects 145
History for Attribute-Based Access Control 145

PART III **Network Address Translation** 147

CHAPTER 10 **Network Address Translation (NAT)** 149
Why use NAT? 149
NAT basics 150
NAT terminology 150
NAT types 150
Network object NAT and Twice NAT 151
Network object NAT 151
Twice NAT 151
Comparing Network object NAT and Twice NAT 152
NAT rule order 152
NAT interfaces 154
Guidelines for NAT 155
Firewall mode guidelines for NAT 155
IPv6 NAT guidelines 155
IPv6 NAT best practices 156
Additional guidelines for NAT 156
Network Object NAT Guidelines for Mapped Address Objects 159
Twice NAT Guidelines for Real and Mapped Address Objects 160
Twice NAT Guidelines for Service Objects for Real and Mapped Ports 161
Dynamic NAT 162
About dynamic NAT 162
Dynamic NAT disadvantages and advantages 163
Configure Dynamic Network Object NAT 163
Configure Dynamic Twice NAT 165
Dynamic PAT 170

About dynamic PAT 170
Dynamic PAT disadvantages and advantages 171
PAT pool object guidelines 171
Configure Dynamic Network Object PAT (Hide) 172
Configure Dynamic Network Object PAT Using a PAT Pool 174
Configure Dynamic Twice PAT (Hide) 176
Configure Dynamic Twice PAT Using a PAT Pool 181
Configure PAT with Port Block Allocation 186
Configure Per-Session PAT or Multi-Session PAT (Version 9.0(1) and Higher) 188
Static NAT 189
About static NAT 189
Static NAT with port translation 190
One-to-many static NAT 191
Other mapping scenarios (not recommended) 192
Configure Static Network Object NAT or Static NAT-with-Port-Translation 193
Configure Static Twice NAT or Static NAT-with-Port-Translation 196
Identity NAT 201
Configure Identity Network Object NAT 201
Configure Identity Twice NAT 203
Monitoring NAT 208
History for NAT 208

CHAPTER 11 **NAT Examples and Reference** 215
Examples for Network Object NAT 215
Providing Access to an Inside Web Server (Static NAT) 215
NAT for Inside Hosts (Dynamic NAT) and NAT for an Outside Web Server (Static NAT) 218
Inside Load Balancer with Multiple Mapped Addresses (Static NAT, One-to-Many) 222
Single Address for FTP, HTTP, and SMTP (Static NAT-with-Port-Translation) 224
Examples for Twice NAT 228
Different Translation Depending on the Destination (Dynamic Twice PAT) 228
Different Translation Depending on the Destination Address and Port (Dynamic PAT) 234
NAT in routed and transparent mode 241
NAT in routed mode 241
NAT in transparent mode or within a bridge group 242

Routing NAT Packets 243
Mapped Addresses and Routing 244
Addresses on the same network as the mapped interface 244
Addresses on a unique network 244
The same address as the real address (identity NAT) 244
Transparent Mode Routing Requirements for Remote Networks 246
Determining the Egress Interface 246
NAT for VPN 247
NAT and Remote Access VPN 247
NAT and Site-to-Site VPN 248
NAT and VPN Management Access 251
Troubleshooting NAT and VPN 252
Translating IPv6 networks 252
NAT64/46: translating IPv6 addresses to IPv4 253
NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet 253
NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation 255
NAT66: translating IPv6 addresses to different IPv6 addresses 259
NAT66 Example, Static Translation between Networks 259
NAT66 Example, Simple IPv6 Interface PAT 260
Rewriting DNS queries and responses using NAT 262
DNS Reply Modification, DNS Server on Outside 263
DNS Reply Modification, DNS Server, Host, and Server on Separate Networks 265
DNS Reply Modification, DNS Server on Host Network 266
DNS64 Reply Modification 268
PTR Modification, DNS Server on Host Network 273

CHAPTER 12 **Mapping Address and Port (MAP) 275**
About Mapping Address and Port (MAP) 275
About Mapping Address and Port Translation (MAP-T) 275
Guidelines for Mapping Address and Port (MAP) 276
Configure MAP-T Domains 278
Monitoring MAP 279
Verifying the MAP Domain Configuration 279
Monitoring MAP Syslog Messages 279

History for MAP 280

PART IV **Service Policies and Application Inspection** 281

CHAPTER 13 **Service Policy** 283

About Service Policies 283

The Components of a Service Policy 283

Features Configured with Service Policies 285

Feature Directionality 286

Feature Matching Within a Service Policy 287

Order in Which Multiple Feature Actions are Applied 287

Incompatibility of Certain Feature Actions 288

Feature Matching for Multiple Service Policies 288

Guidelines for Service Policies 289

Defaults for Service Policies 290

Default Service Policy Configuration 290

Default Class Maps (Traffic Classes) 291

Configure Service Policies 291

Add a Service Policy Rule for Through Traffic 291

Add a Service Policy Rule for Management Traffic 294

Manage the Order of Service Policy Rules 296

History for Service Policies 297

CHAPTER 14 **Getting Started with Application Layer Protocol Inspection** 299

Application Layer Protocol Inspection 299

When to Use Application Protocol Inspection 299

Inspection Policy Maps 300

Replacing an In-Use Inspection Policy Map 300

How Multiple Traffic Classes are Handled 300

Guidelines for Application Inspection 301

Defaults for Application Inspection 302

Default Inspections and NAT Limitations 302

Default Inspection Policy Maps 306

Configure Application Layer Protocol Inspection 306

Configure Regular Expressions 310
Create a Regular Expression 310
Create a Regular Expression Class Map 314
Monitoring Inspection Policies 314
History for Application Inspection 315

CHAPTER 15 **Inspection of Basic Internet Protocols** 317
DCERPC Inspection 317
DCERPC Overview 318
Configure a DCERPC Inspection Policy Map 318
DNS Inspection 320
Defaults for DNS Inspection 320
Configure DNS Inspection Policy Map 321
FTP Inspection 324
FTP Inspection Overview 324
Strict FTP 324
Configure an FTP Inspection Policy Map 325
HTTP Inspection 328
HTTP Inspection Overview 328
Configure an HTTP Inspection Policy Map 329
ICMP Inspection 332
ICMP Error Inspection 332
ILS Inspection 333
Instant Messaging Inspection 333
IP Options Inspection 335
Defaults for IP Options Inspection 336
Configure an IP Options Inspection Policy Map 336
IPsec Pass Through Inspection 337
IPsec Pass Through Inspection Overview 337
Configure an IPsec Pass Through Inspection Policy Map 338
IPv6 Inspection 339
Defaults for IPv6 Inspection 339
Configure an IPv6 Inspection Policy Map 339
NetBIOS Inspection 340

PPTP Inspection 341
RSH Inspection 341
SMTP and Extended SMTP Inspection 341
SMTP and ESMTP Inspection Overview 342
Defaults for ESMTP Inspection 343
Configure an ESMTP Inspection Policy Map 343
SNMP Inspection 345
SQL*Net Inspection 346
Sun RPC Inspection 346
Sun RPC Inspection Overview 347
Manage Sun RPC Services 347
TFTP Inspection 348
XDMCP Inspection 348
VXLAN Inspection 349
History for Basic Internet Protocol Inspection 349

CHAPTER 16 **Inspection for Voice and Video Protocols** 351
CTIQBE Inspection 351
Limitations for CTIQBE Inspection 351
H.323 Inspection 352
H.323 Inspection Overview 352
How H.323 Works 352
H.239 Support in H.245 Messages 353
Limitations for H.323 Inspection 354
Configure H.323 Inspection Policy Map 354
MGCP Inspection 356
MGCP Inspection Overview 356
Configure an MGCP Inspection Policy Map 358
RTSP Inspection 359
RTSP Inspection Overview 359
RealPlayer Configuration Requirements 359
Limitations for RSTP Inspection 359
Configure RTSP Inspection Policy Map 360
SIP Inspection 361

SIP Inspection Overview 362
Limitations for SIP Inspection 362
Default SIP Inspection 363
Configure SIP Inspection Policy Map 364
Skinny (SCCP) Inspection 366
SCCP Inspection Overview 366
Supporting Cisco IP Phones 367
Limitations for SCCP Inspection 367
Default SCCP Inspection 367
Configure a Skinny (SCCP) Inspection Policy Map 368
STUN Inspection 369
History for Voice and Video Protocol Inspection 370

CHAPTER 17 **Inspection for Mobile Networks 373**
Mobile Network Inspection Overview 373
GTP Inspection Overview 373
Tracking Location Changes for Mobile Stations 374
GTP Inspection Limitations 374
Stream Control Transmission Protocol (SCTP) Inspection and Access Control 374
SCTP Stateful Inspection 375
SCTP Access Control 376
SCTP NAT 376
SCTP Application Layer Inspection 376
SCTP Limitations 377
Diameter Inspection 377
M3UA Inspection 378
M3UA Protocol Conformance 378
M3UA Inspection Limitations 379
RADIUS Accounting Inspection Overview 379
Licensing for Mobile Network Protocol Inspection 380
Defaults for GTP Inspection 380
Configure Mobile Network Inspection 381
Configure a GTP Inspection Policy Map 381
Configure an SCTP Inspection Policy Map 385

Configure a Diameter Inspection Policy Map 386
Create a Custom Diameter Attribute-Value Pair (AVP) 389
Inspecting Encrypted Diameter Sessions 390
Configure Server Trust Relationship with Diameter Clients 391
Configure Full TLS Proxy with Static Client Certificate for Diameter Inspection 392
Configure Full TLS Proxy with Local Dynamic Certificates for Diameter Inspection 393
Configure TLS Proxy with TLS Offload for Diameter Inspection 395
Configure an M3UA Inspection Policy Map 397
Configure the Mobile Network Inspection Service Policy 399
Configure RADIUS Accounting Inspection 400
Configure a RADIUS Accounting Inspection Policy Map 401
Configure the RADIUS Accounting Inspection Service Policy 402
Monitoring Mobile Network Inspection 402
Monitoring GTP Inspection 402
Monitoring SCTP 404
Monitoring Diameter 405
Monitoring M3UA 405
History for Mobile Network Inspection 406

PART V **Connection Management and Threat Detection** 411

CHAPTER 18 **Connection Settings** 413
Connection settings 413
Configure Connection Settings 414
Configure Global Timeouts 415
Protect Servers from a SYN Flood DoS Attack (TCP Intercept) 417
Customize Abnormal TCP Packet Handling (TCP Maps, TCP Normalizer) 419
Bypass TCP state checks for asymmetrical routing (TCP state bypass) 422
The asymmetrical routing problem 422
Guidelines and limitations for TCP state bypass 423
Configure TCP State Bypass 424
Disable TCP Sequence Randomization 424
Offload Large Flows 425
Restrictions for flow offload implementation 426

Configure Flow Offload 427
Configure Connection Settings for Specific Traffic Classes (All Services) 428
Configure TCP Options 431
Monitoring Connections 432
History for Connection Settings 433

CHAPTER 19 **Quality of Service** 437
About QoS 437
Supported QoS Features 437
What is a Token Bucket? 438
Policing 438
Priority Queuing 438
How QoS Features Interact 439
DSCP (DiffServ) Preservation 439
Guidelines for QoS 439
Configure QoS 439
Determine the Queue and TX Ring Limits for a Priority Queue 440
Queue Limit Worksheet 440
TX Ring Limit Worksheet 440
Configure the Priority Queue for an Interface 441
Configure a Service Rule for Priority Queuing and Policing 442
Monitor QoS 443
QoS Police Statistics 443
QoS Priority Statistics 444
QoS Priority Queue Statistics 444
History for QoS 445

CHAPTER 20 **Threat Detection** 447
Detecting Threats 447
Basic Threat Detection Statistics 448
Advanced Threat Detection Statistics 448
Scanning Threat Detection 449
Guidelines for Threat Detection 449
Defaults for Threat Detection 450

Configure Threat Detection 451
Configure Basic Threat Detection Statistics 451
Configure Advanced Threat Detection Statistics 451
Configure Scanning Threat Detection 452
Monitoring Threat Detection 453
Monitoring Basic Threat Detection Statistics 453
Monitoring Advanced Threat Detection Statistics 453
History for Threat Detection 454



About This Guide

The following topics explain how to use this guide.

- Document Objectives, on page xix
- Related Documentation, on page xix
- Document Conventions, on page xix
- Communications, Services, and Additional Information, on page xxi

Document Objectives

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

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

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

Related Documentation

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

Document Conventions

This document adheres to the following text, display, and alert conventions.

Text Conventions

Convention	Indication
boldface	Commands, keywords, button labels, field names, and user-entered text appear in **boldface**. For menu-based commands, the full path to the command is shown.
italic	Variables, for which you supply values, are presented in an *italic* typeface. Italic type is also used for document titles, and for general emphasis.
`monospace`	Terminal sessions and information that the system displays appear in `monospace` type.
{x | y | z}	Required alternative keywords are grouped in braces and separated by vertical bars.
[ ]	Elements in square brackets are optional.
[x | y | z]	Optional alternative keywords are grouped in square brackets and separated by vertical bars.
[ ]	Default responses to system prompts are also in square brackets.
< >	Non-printing characters such as passwords are in angle brackets.
!, #	An exclamation point (!) or a number sign (#) at the beginning of a line of code indicates a comment line.

Reader Alerts

This document uses the following for reader alerts:



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



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



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



Timesaver Means *the described action saves time.* You can save time by performing the action described in the paragraph.



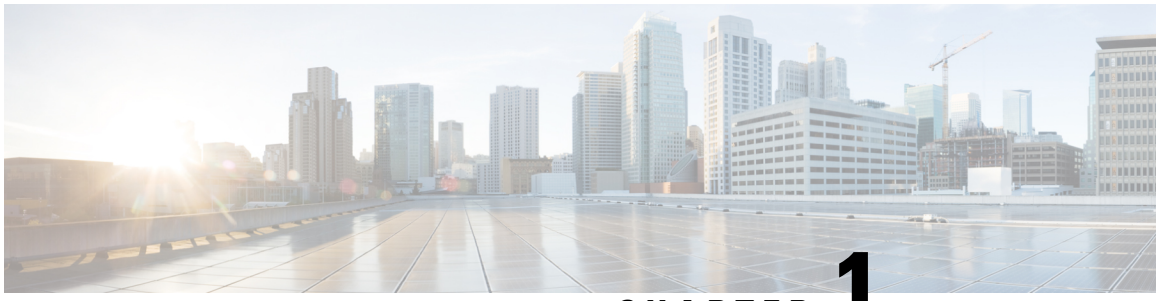
Warning **Means *reader be warned.* In this situation, you might perform an action that could result in bodily injury.**

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business results you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 1

Introduction to ASA-Firewall Services

Firewall services are those ASA features that are focused on controlling access to the network, including services that block traffic and services that enable traffic flow between internal and external networks. These services include those that protect the network against threats, such as Denial of Service (DoS) and other attacks.

The following topics provide an overview of firewall services.

- How to Implement Firewall Services, on page 1
- Basic Access Control, on page 2
- Application Filtering, on page 2
- URL Filtering, on page 2
- Threat Protection, on page 3
- Firewall Services for Virtual Environments, on page 4
- Network Address Translation, on page 4
- Application Inspection, on page 5
- Use Case: Expose a Server to the Public, on page 5

How to Implement Firewall Services

The following procedure provides a general sequence for implementing firewall services. However, each step is optional, needed only if you want to provide the service to your network.

Before you begin

Configure the ASA according to the general operations configuration guide, including at minimum basic settings, interface configuration, routing, and management access.

Procedure

Step 1 Implement access control for the network. See Basic Access Control, on page 2.

Step 2 Implement application filtering. See Application Filtering, on page 2.

Step 3 Implement URL filtering. See URL Filtering, on page 2.

Step 4 Implement threat protection. See Threat Protection, on page 3.

Step 5 Implement firewall services that are tailored to virtual environments. See Firewall Services for Virtual Environments, on page 4.

Step 6 Implement Network Address Translation (NAT). See Network Address Translation, on page 4.

Step 7 Implement application inspection if the default settings are insufficient for your network. See Application Inspection, on page 5.

Basic Access Control

Access rules, applied per interface or globally, are your first line of defense. You can drop, upon entry, specific types of traffic, or traffic from (or to) specific hosts or networks. By default, the ASA allows traffic to flow freely from an inside network (higher security level) to an outside network (lower security level).

You can apply an access rule to limit traffic from inside to outside, or allow traffic from outside to inside.

Basic access rules control traffic using a "5-tuple" of source address and port, destination address and port, and protocol. See Access Rules, on page 11 and Access Control Lists, on page 41.

You can augment your rules by making them identity aware. This lets you configure rules based on user identity or group membership. To implement identity control, do any combination of the following:

- Install Cisco Identity Services Engine (ISE) on a separate server to implement Cisco Trustsec. You can then add security group criteria to your access rules. See ASA and Cisco TrustSec, on page 73.
- Install the ASA FirePOWER module on the ASA and implement identity policies in the module. The identity-aware access policies in ASA FirePOWER would apply to any traffic that you redirect to the module. See ASA FirePOWER Module, on page 97.

Application Filtering

The wide-spread use of web-based applications means that a lot of traffic runs over the HTTP or HTTPS protocols. With traditional 5-tuple access rules, you either allow or disallow all HTTP/HTTPS traffic. You might require more granular control of web traffic.

You can install a module on the ASA to provide application filtering to selectively allow HTTP or other traffic based on the application being used. Thus, you do not have to make a blanket permit for HTTP. You can look inside the traffic and prevent applications that are unacceptable for your network (for example, inappropriate file sharing). When you add a module for application filtering, do not configure HTTP inspection on the ASA.

To implement application filtering, install the ASA FirePOWER module on the ASA and use application filtering criteria in your ASA FirePOWER access rules. These policies apply to any traffic that you redirect to the module. See ASA FirePOWER Module, on page 97.

URL Filtering

URL filtering denies or allows traffic based on the URL of the destination site.

The purpose of URL filtering is primarily to completely block or allow access to a web site. Although you can target individual pages, you typically specify a host name (such as www.example.com) or a URL category, which defines a list of host names that provide a particular type of service (such as Gambling).

When trying to decide whether to use URL filtering or application filtering for HTTP/HTTPS traffic, consider whether your intention is to create a policy that applies to all traffic directed at a web site. If your intention is to treat all such traffic the same way (denying it or allowing it), use URL filtering. If your intention is to selectively block or allow traffic to the site, use application filtering.

To implement URL filtering, do one of the following:

- Install the ASA FirePOWER module on the ASA and use URL filtering criteria in your ASA FirePOWER access rules. These policies apply to any traffic that you redirect to the module. See ASA FirePOWER Module, on page 97.
- Subscribe to the Cisco Umbrella service, where you configure the Enterprise Security policy to block malicious sites based on the fully-qualified domain name (FQDN). For FQDNs that are considered suspicious, you can redirect user connections to the Cisco Umbrella intelligent proxy, which performs URL filtering. The Umbrella service works by handling users' DNS lookup requests, returning the IP address for a block page or the IP address of the intelligent proxy. The service returns the real IP address for an FQDN for allowed domains. See Cisco Umbrella, on page 125.

Threat Protection

You can implement a number of measures to protect against scanning, denial of service (DoS), and other attacks. A number of ASA features help protect against attacks by applying connection limits and dropping abnormal TCP packets. Some features are automatic, others are configurable but have defaults appropriate in most cases, while others are completely optional and you must configure them if you want them.

Following are the threat protection services available with the ASA.

- IP packet fragmentation protection—The ASA performs full reassembly of all ICMP error messages and virtual reassembly of the remaining IP fragments that are routed through the ASA, and drops fragments that fail the security check. No configuration is necessary.
- Connection limits, TCP normalization, and other connection-related features—Configure connection-related services such as TCP and UDP connection limits and timeouts, TCP sequence number randomization, TCP normalization, and TCP state bypass. TCP normalization is designed to drop packets that do not appear normal. See Connection Settings, on page 413.

  For example, you can limit TCP and UDP connections and embryonic connections (a connection request that has not finished the necessary handshake between source and destination). Limiting the number of connections and embryonic connections protects you from a DoS attack. The ASA uses the embryonic limit to trigger TCP Intercept, which protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets.
- Threat detection—Implement threat detection on the ASA to collect statistics to help identify attacks. Basic threat detection is enabled by default, but you can implement advanced statistics and scanning threat detection. You can shun hosts that are identified as a scanning threat. See Threat Detection, on page 447.
- Next-Generation IPS—Install the ASA FirePOWER module on the ASA and implement Next Generation IPS intrusion rules in your ASA FirePOWER. These policies would apply to any traffic that you redirect to ASA FirePOWER. See ASA FirePOWER Module, on page 97.

Firewall Services for Virtual Environments

Virtual environments deploy servers as virtual machines, for example, in VMware ESXi. The firewalls in a virtual environment can be traditional hardware devices, or they can also be virtual machine firewalls, such as the ASAv.

Traditional and next-generation firewall services apply to virtual environments in the same way that they apply to environments that do not use virtual machine servers. However, virtual environments can provide additional challenges, because it is easy to create and tear down servers.

Additionally, traffic between servers within the data center might require as much protection as traffic between the data center and external users. For example, if an attacker gains control of a server within the data center, that could open up attacks on other servers in the data center.

Firewall services for virtual environments add capabilities to apply firewall protection specifically to virtual machines. Following are the firewall services available for virtual environments:

- Attribute-based access control—You can configure network objects to match traffic based on attributes, and use those objects in access control rules. This lets you decouple firewall rules from network topology. For example, you can allow all hosts with the Engineering attribute to access hosts with the Lab Server attribute. You could then add/remove hosts with these attributes and the firewall policy would be applied automatically without the need for updating access rules. For more information, see Attribute-Based Access Control, on page 139.

Network Address Translation

One of the main functions of Network Address Translation (NAT) is to enable private IP networks to connect to the Internet. NAT replaces a private IP address with a public IP address, translating the private addresses in the internal private network into legal, routable addresses that can be used on the public Internet. In this way, NAT conserves public addresses because you can advertise at a minimum only one public address for the entire network to the outside world.

Other functions of NAT include:

- Security—Keeping internal IP addresses hidden discourages direct attacks.
- IP routing solutions—Overlapping IP addresses are not a problem when you use NAT.
- Flexibility—You can change internal IP addressing schemes without affecting the public addresses available externally; for example, for a server accessible to the Internet, you can maintain a fixed IP address for Internet use, but internally, you can change the server address.
- Translating between IPv4 and IPv6 (Routed mode only)—If you want to connect an IPv6 network to an IPv4 network, NAT lets you translate between the two types of addresses.

NAT is not required. If you do not configure NAT for a given set of traffic, that traffic will not be translated, but will have all of the security policies applied as normal.

See:

- Network Address Translation (NAT), on page 149
- NAT Examples and Reference, on page 215

Application Inspection

Application inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection, to open the required pinholes and to apply network address translation (NAT).

The default ASA policy already applies inspection globally for many popular protocols, such as DNS, FTP, SIP, ESMTP, TFTP, and others. The default inspections might be all you require for your network.

However, you might need to enable inspection for other protocols, or fine-tune an inspection. Many inspections include detailed options that let you control packets based on their contents. If you know a protocol well, you can apply fine-grained control on that traffic.

You use service policies to configure application inspection. You can configure a global service policy, or apply a service policy to each interface, or both.

See:

- Service Policy, on page 283
- Getting Started with Application Layer Protocol Inspection, on page 299
- Inspection of Basic Internet Protocols, on page 317
- Inspection for Voice and Video Protocols, on page 351
- Inspection for Mobile Networks, on page 373.

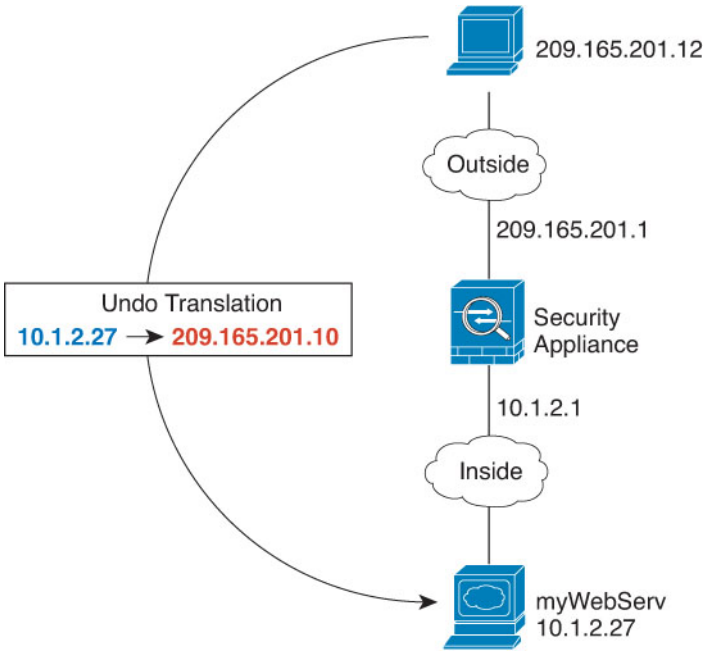
Use Case: Expose a Server to the Public

You can make certain application services on a server available to the public. For example, you could expose a web server, so that users can connect to the web pages but not make any other connections to the server.

To expose a server to the public, you typically need to create access rules that allow the connection and NAT rules to translate between the server's internal IP address and an external address that the public can use. In addition, you can use port address translation (PAT) to map an internal port to an external port, if you do not want the externally exposed service to use the same port as the internal server. For example, if the internal web server is not running on TCP/80, you can map it to TCP/80 to make connections easier for external users.

The following example makes a web server on the inside private network available for public access.

Figure 1: Static NAT for an Inside Web Server



ASDM includes a short cut for configuring the required access and NAT rules, to simplify the process of exposing a service on an internal server to the public.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Public Servers**.

Step 2 Click **Add**.

Step 3 Define the private and public characteristics of the service you are exposing.

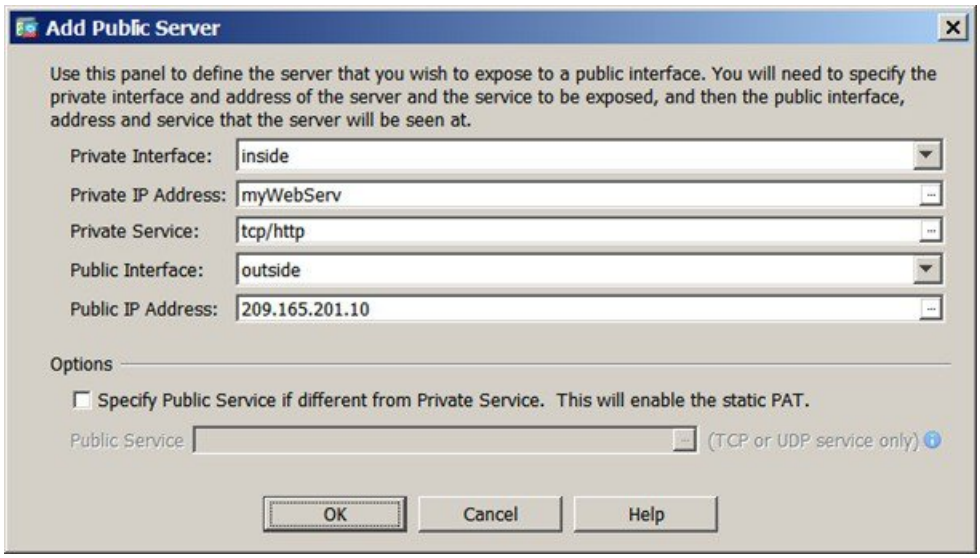
- **Private Interface**—The interface to which the real server is connected. In this example, "inside."
- **Private IP Address**—The host network object that defines the real IPv4 address of the server. You cannot specify an IPv6 address. If you do not already have an object containing the address, create one by clicking the "..." button and then clicking **Add**. In this example, the object name would be MyWebServ, and it would contain the 10.1.2.27 host address.
- **Private Service**—The actual service that is running on the real server. You can use a pre-defined service or service object. You can also use a service object group unless you also specify a public service to which you are mapping the private service.

  You can expose multiple services; however, if you specify a public service, all ports are mapped to the same public port.

  In this example, the port is **tcp/http**.
- **Public Interface**—The interface through which outside users can access the real server. In this example, "outside."

- **Public Address**—The IPv4 address that is seen by outside users. You can specify the address directly or use a host network object. In this example, the outside address is 209.165.201.10.
- **Specify Public Service if different from private service**, **Public Service**—The service that is running on the translated address. Specify the public service only if it differs from the private service. For example, if the private web server runs on TCP/80, and you want to use the same port for external users, there is no need to specify the public service. You must use a pre-defined TCP or UDP service if you specify a public service. This example does not use port translation, so do not select this option.

The following shows how the dialog box should look.



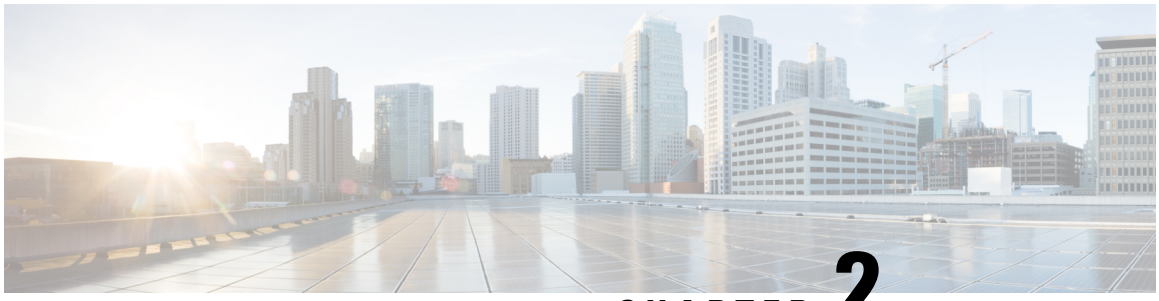
Step 4 Click **OK**, then click **Apply**.



PART I

Access Control

- Access Rules, on page 11
- Objects for Access Control, on page 31
- Access Control Lists, on page 41
- Identity Firewall, on page 57
- ASA and Cisco TrustSec, on page 73
- ASA FirePOWER Module, on page 97
- Cisco Umbrella, on page 125



CHAPTER 2

Access Rules

This chapter describes how to control network access through or to the ASA using access rules. You use access rules to control network access in both routed and transparent firewall modes. In transparent mode, you can use both access rules (for Layer 3 traffic) and EtherType rules (for Layer 2 traffic).

Note To access the ASA interface for management access, you do not also need an access rule allowing the host IP address. You only need to configure management access according to the general operations configuration guide.

- Controlling Network Access, on page 11
- Licensing for Access Rules, on page 17
- Guidelines for Access Control, on page 17
- Configure Access Control, on page 19
- Monitoring Access Rules, on page 27
- History for Access Rules, on page 28

Controlling Network Access

Access rules determine which traffic is allowed through the ASA. There are several different layers of rules that work together to implement your access control policy:

- Extended access rules (Layer 3+ traffic) assigned to interfaces—You can apply separate rule sets (ACLs) in the inbound and outbound directions. An extended access rule permits or denies traffic based on the source and destination traffic criteria.
- Extended access rules (Layer 3+ traffic) assigned to Bridge Virtual Interfaces (BVI; routed mode)—If you name a BVI, you can apply separate rule sets in the inbound and outbound direction, and you can also apply rule sets to the bridge group member interfaces. When both the BVI and member interface have access rules, the order of processing depends on direction. Inbound, the member access rules are evaluated first, then the BVI access rules. Outbound, the BVI rules are considered first, then the member interface rules.
- Extended access rules assigned globally—You can create a single global rule set, which serves as your default access control. The global rules are applied after interface rules.

- Management access rules (Layer 3+ traffic)—You can apply a single rule set to cover traffic directed at an interface, which would typically be management traffic. In the CLI, these are “control plane” access groups. For ICMP traffic directed at the device, you can alternatively configure ICMP rules.
- EtherType rules (Layer 2 traffic) assigned to interfaces (bridge group member interfaces only)—You can apply separate rule sets in the inbound and outbound directions. EtherType rules control network access for non-IP traffic. An EtherType rule permits or denies traffic based on the EtherType. You can also apply extended access rules to bridge group member interfaces to control Layer 3+ traffic.

General Information About Rules

The following topics provide general information about access rules and EtherType rules.

Interface Access Rules and Global Access Rules

You can apply an access rule to a specific interface, or you can apply an access rule globally to all interfaces. You can configure global access rules in conjunction with interface access rules, in which case, the specific inbound interface access rules are always processed before the general global access rules. Global access rules apply only to inbound traffic.

Inbound and Outbound Rules

You can configure access rules based on the direction of traffic:

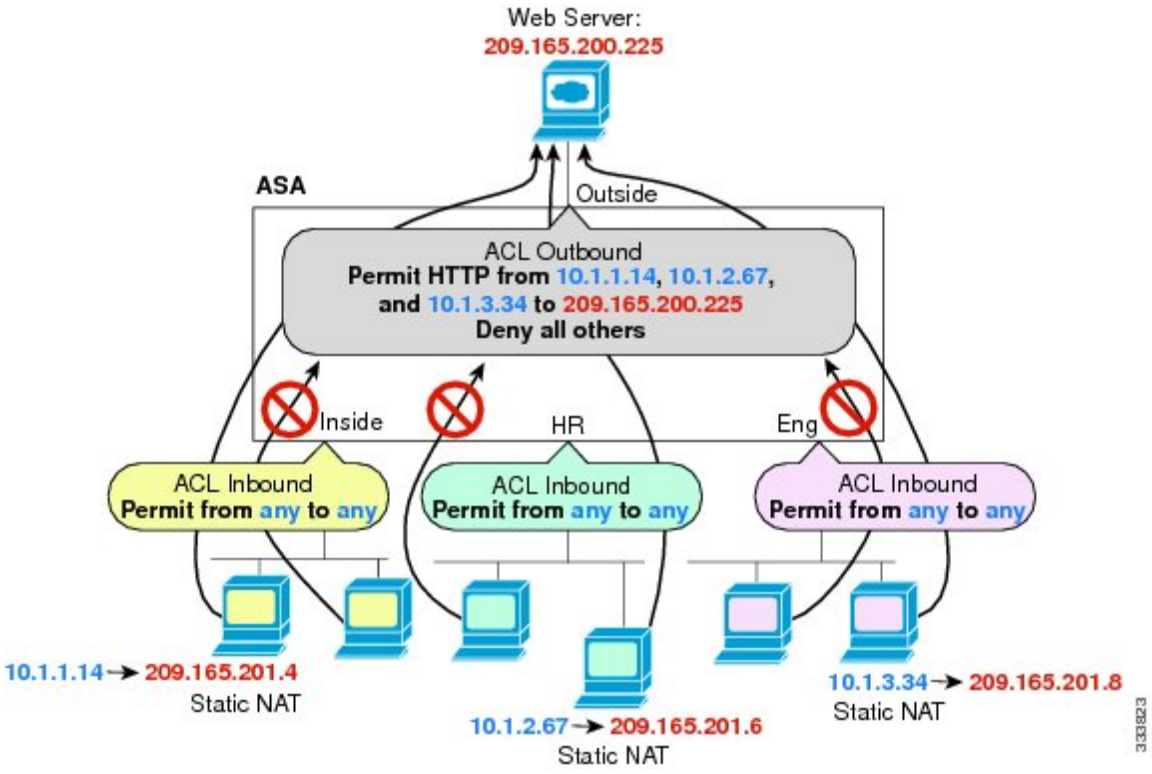
- Inbound—Inbound access rules apply to traffic as it enters an interface. Global and management access rules are always inbound.
- Outbound—Outbound rules apply to traffic as it exits an interface.



Note “Inbound” and “outbound” refer to the application of an ACL on an interface, either to traffic entering the ASA on an interface or traffic exiting the ASA on an interface. These terms do not refer to the movement of traffic from a lower security interface to a higher security interface, commonly known as inbound, or from a higher to lower interface, commonly known as outbound.

An outbound ACL 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 ACLs to restrict access, you can create a single outbound ACL that allows only the specified hosts. (See the following figure.) The outbound ACL prevents any other hosts from reaching the outside network.

Figure 2: Outbound ACL



Rule Order

The order of rules is important. When the ASA decides whether to forward or drop a packet, the ASA tests the packet against each rule in the order in which the rules are listed in the applied ACL. After a match is found, no more rules are checked. For example, if you create an access rule at the beginning that explicitly permits all traffic for an interface, no further rules are ever checked.

Implicit Permits

Unicast IPv4 and IPv6 traffic from a higher security interface to a lower security interface is allowed through by default. This includes traffic between standard routed interfaces and Bridge Virtual Interfaces (BVI) in routed mode.

For bridge group member interfaces, this implicit permit from a higher to a lower security interface applies to interfaces within the same bridge group only. There are no implicit permits between a bridge group member interface and a routed interface or a member of a different bridge group.

Bridge group member interfaces (routed or transparent mode) also allow the following by default:

- ARPs in both directions. (You can control ARP traffic using ARP inspection, but you cannot control it by access rule.)
- BPDUs in both directions. (You can control these using Ethertype rules.)

For other traffic, you need to use either an extended access rule (IPv4 and IPv6) or an EtherType rule (non-IP).

Implicit Deny

ACLs have an implicit deny at the end of the list, so unless you explicitly permit it, traffic cannot pass. For example, if you want to allow all users to access a network through the ASA except for particular addresses, then you need to deny the particular addresses and then permit all others.

For management (control plane) ACLs, which control to-the-box traffic, there is no implicit deny at the end of a set of management rules for an interface. Instead, any connection that does not match a management access rule is then evaluated by regular access control rules.

For EtherType ACLs, the implicit deny at the end of the ACL does not affect IP traffic or ARPs; for example, if you allow EtherType 8037, the implicit deny at the end of the ACL does not now block any IP traffic that you previously allowed with an extended ACL (or implicitly allowed from a high security interface to a low security interface). However, if you explicitly deny all traffic with an EtherType rule, then IP and ARP traffic is denied; only physical protocol traffic, such as auto-negotiation, is still allowed.

If you configure a global access rule, then the implicit deny comes *after* the global rule is processed. See the following order of operations:

1. Interface access rule.
2. For bridge group member interfaces, the Bridge Virtual Interface (BVI) access rule.
3. Global access rule.
4. Implicit deny.

NAT and Access Rules

Access rules always use the real IP addresses when determining an access rule match, even if you configure NAT. For example, if you configure NAT for an inside server, 10.1.1.5, so that it has a publicly routable IP address on the outside, 209.165.201.5, then the access rule to allow the outside traffic to access the inside server needs to reference the server's real IP address (10.1.1.5), and not the mapped address (209.165.201.5).

Same Security Level Interfaces and Access Rules

Each interface has a security level, and security level checking is performed before access rules are considered. Thus, even if you allow a connection in an access rule, it can be blocked due to same-security-level checking at the interface level. You might want to ensure that your configuration allows same-security-level connections so that your access rules are always considered for permit/deny decisions.

- Connections between the same security level ingress and egress interfaces are subject to the same-security-traffic inter-interface check.

  To allow these connections, enter the **same-security-traffic permit inter-interface** command.

  To allow these connections, choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, then select the **Enable traffic between two or more interfaces which are configured with the same security levels** option.

- Connections with the same ingress and egress interfaces are subject to the same-security-traffic intra-interface check.

  To allow these connections, enter the **same-security-traffic permit intra-interface** command.

  To allow these connections, choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, then select the **Enable traffic between two or more hosts connected to the same interface** option.

Extended Access Rules

This section describes information about extended access rules.

Extended Access Rules for Returning Traffic

For TCP, UDP, and SCTP connections for both routed and transparent mode, you do not need an access rule to allow returning traffic because the ASA allows all returning traffic for established, bidirectional connections.

For connectionless protocols such as ICMP, however, the ASA establishes unidirectional sessions, so you either need access rules to allow ICMP in both directions (by applying ACLs 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. For example, to control ping, specify **echo-reply** (**0**) (ASA to host) or **echo** (**8**) (host to ASA).

Allowing Broadcast and Multicast Traffic

In routed firewall mode, broadcast and multicast traffic is blocked even if you allow it in an access rule, including unsupported dynamic routing protocols and DHCP. You must configure the dynamic routing protocols or DHCP relay to allow this traffic.

For interfaces that are members of the same bridge group in transparent or routed firewall mode, you can allow any IP traffic through using access rules.



Note Because these special types of traffic are connectionless, you need to apply an access rule to both the inbound and outbound interfaces, so returning traffic is allowed through.

The following table lists common traffic types that you can allow using access rules between interfaces that are members of the same bridge group.

Table 1: Special Traffic for Access Rules between Members of the Same Bridge Group

Traffic Type	Protocol or Port	Notes
DHCP	UDP ports 67 and 68	If you enable the DHCP server, then the ASA does not pass DHCP packets.
EIGRP	Protocol 88	—
OSPF	Protocol 89	—
Multicast streams	The UDP ports vary depending on the application.	Multicast streams are always destined to a Class D address (224.0.0.0 to 239.x.x.x).
RIP (v1 or v2)	UDP port 520	—

Management Access Rules

You can configure access rules that control management traffic destined to the ASA. Access control rules for to-the-box management traffic (such as HTTP, Telnet, and SSH connections to an interface) have higher

precedence than a management access rule . Therefore, such permitted management traffic will be allowed to come in even if explicitly denied by the to-the-box ACL.

Unlike regular access rules, there is no implicit deny at the end of a set of management rules for an interface. Instead, any connection that does not match a management access rule is then evaluated by regular access control rules.

Alternatively, you can use ICMP rules to control ICMP traffic to the device. Use regular extended access rules to control ICMP traffic through the device.

EtherType Rules

This section describes EtherType rules.

Supported EtherTypes and Other Traffic

An EtherType rule controls the following:

- EtherType identified by a 16-bit hexadecimal number, including common types IPX and MPLS unicast or multicast.
- Ethernet V2 frames.
- BPDUs, which are permitted by default. BPDUs are SNAP-encapsulated, and the ASA is designed to specifically handle BPDUs.
- Trunk port (Cisco proprietary) BPDUs. Trunk BPDUs have VLAN information inside the payload, so the ASA modifies the payload with the outgoing VLAN if you allow BPDUs.
- Intermediate System to Intermediate System (IS-IS).
- The IEEE 802.2 Logical Link Control packet. You can control access based on the Destination Service Access Point address.

The following types of traffic are not supported:

- 802.3-formatted frames—These frames are not handled by the rule because they use a length field as opposed to a type field.

EtherType Rules for Returning Traffic

Because EtherTypes are connectionless, you need to apply the rule to both interfaces if you want traffic to pass in both directions.

Allowing MPLS

If you allow MPLS, ensure that Label Distribution Protocol and Tag Distribution Protocol TCP connections are established through the ASA by configuring both MPLS routers connected to the ASA to use the IP address on the ASA interface as the router-id for LDP or TDP sessions. (LDP and TDP allow MPLS routers to negotiate the labels (addresses) used to forward packets.)

On Cisco IOS routers, enter the appropriate command for your protocol, LDP or TDP. The *interface* is the interface connected to the ASA.

mpls ldp router-id *interface* **force**

Or

tag-switching tdp router-id *interface* **force**

Licensing for Access Rules

Access control rules do not require a special license.

However, to use **sctp** as the protocol in a rule, you must have a Carrier license.

Guidelines for Access Control

IPv6 Guidelines

Supports IPv6. (9.0 and later) The source and destination addresses can include any mix of IPv4 and IPv6 addresses. For pre-9.0 versions, you must create a separate IPv6 access rule.

Per-User ACL Guidelines

- The per-user ACL uses the value in the **timeout uauth** command, but it can be overridden by the AAA per-user session timeout value.
- If traffic is denied because of a per-user ACL, syslog message 109025 is logged. If traffic is permitted, no syslog message is generated. The **log** option in the per-user ACL has no effect.

Additional Guidelines and Limitations

- Over time, your list of access rules can grow to include many obsolete rules. Eventually, the ACLs for the access groups can become so large that they impact overall system performance. If you find that the system is having issues sending syslog messages, communicating for failover synchronization, establishing and maintaining SSH/HTTPS management access connections, and so forth, you might need to prune your access rules. In general, you should actively maintain your rule lists to remove obsolete rules, rules that are never hit, FQDN objects that can no longer be resolved, and so forth. Also consider implementing object group search.
- You can reduce the memory required to search access rules by enabling object group search, but this is at the expense rule of lookup performance and increased CPU utilization. When enabled, object group search does not expand network or service objects, but instead searches access rules for matches based on those group definitions. You can set this option by clicking the **Advanced** button below the access rule table.

  You can use the **object-group-search threshold** command to enable a threshold to help prevent performance degradation. When operating with a threshold, for each connection, both the source and destination IP addresses are matched against network objects. If the number of objects matched by the source address times the number matched by the destination address exceeds 10,000, the connection is dropped. Configure your rules to prevent an excessive number of matches.



Note Object group search works with network and service objects only. It does not work with security group or user objects. Do not enable this feature if the ACLs include security groups. The result can be inactive ACLs or other unexpected behavior.

- You can improve system performance and reliability by using the transactional commit model for access groups. However, transactional commit is not recommended if you use FQDN objects for hostnames whose resolution can frequently change, as the access group compilation might never completely resolve. See the basic settings chapter in the general operations configuration guide for more information. The option is under **Configurations > Device Management > Advanced > Rule Engine**.
- In ASDM, rule descriptions are based on the access list remarks that come before the rule in the ACL; for new rules you create in ASDM, any descriptions are also configured as remarks before the related rule. However, the packet tracer in ASDM matches the remark that is configured after the matching rule in the CLI.
- If you enter more than one item in source or destination address, or source or destination service, ASDM automatically creates an object group for them with the prefix DM_INLINE. These objects are automatically expanded to their component parts in the rule table view, but you can see the object names if you deselect the **Auto-expand network and service objects with specified prefix** rule table preference in **Tools** > **Preferences**.
- Normally, you cannot reference an object or object group that does not exist in an ACL or object group, or delete one that is currently referenced. You also cannot reference an ACL that does not exist in an **access-group** command (to apply access rules). However, you can change this default behavior so that you can "forward reference" objects or ACLs before you create them. Until you create the objects or ACLs, any rules or access groups that reference them are ignored. To enable forward referencing, select the option in the access rules advanced settings; choose **Configuration > Access Rules** and click the **Advanced** button.
- To use fully-qualified domain name (FQDN) network objects as source or destination criteria, you must also configure DNS for the data interfaces.

  Note that controlling access by FQDN is a best-effort mechanism. Consider the following points:

  - Because DNS replies can be spoofed, only use fully trusted internal DNS servers.
  - Some FQDNs, especially for very popular servers, can have hundreds if not thousands of IP addresses, and these can frequently change. Because the system uses cached DNS lookup results, users might get addresses that are not yet in the cache, and their connections will not match the FQDN rule. Rules that use FQDN network objects function effectively only for names that resolve to fewer than 100 addresses.

    We recommend that you do not create network object rules for an FQDN that resolves to more than 100 addresses, as the likelihood of the address in a connection being one that has been resolved and available in the DNS cache on the device is low.
  - For popular FQDNs, different DNS servers can return a different set of IP addresses. Thus, if your users use a different DNS server than the one you configure, FQDN-based access control rules might not apply to all IP addresses for the site that are used by your clients, and you will not get the intended results for your rules.

- Some FQDN DNS entries have very short time to live (TTL) values. This can result in frequent recompliation of the lookup table, which can impact overall system performance.

Configure Access Control

The following topics explain how to configure access control.

Configure Access Rules

To apply an access rule, perform the following steps.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Access Rules**.

The rules are organized by interface and direction, with a separate group for global rules. If you configure management access rules, they are repeated on this page. These groups are equivalent to the extended ACL that is created and assigned to the interface or globally as an access group. These ACLs also appear on the ACL Manager page.

Step 2 Do any of the following:

- To add a new rule, choose **Add > Add Access Rule**.
- To insert a rule at a specific location within a container, select an existing rule and choose **Add > Insert** to add the rule above it, or choose **Add > Insert After**.
- To edit a rule, select it and click **Edit**.

Step 3 Fill in the rule properties. The primary options to select are:

- **Interface**—The interface to which the rule applies. Select Any to create a global rule. For bridge groups in routed mode, you can create access rules for both the Bridge Virtual Interface (BVI) and each bridge group member interface.
- **Action: Permit/Deny**—Whether you are permitting (allowing) the described traffic or are denying (dropping) it.
- **Source/Destination criteria**—A definition of the source (originating address) and destination (target address of the traffic flow). You typically configure IPv4 or IPv6 addresses of hosts or subnets, which you can represent with network or network object groups. You can also specify a user or user group name for the source. Additionally, you can use the Service field to identify the specific type of traffic if you want to focus the rule more narrowly than all IP traffic. If you implement Trustsec, you can use security groups to define source and destination.

For detailed information on all of the available options, see Access Rule Properties, on page 20.

When you are finished defining the rule, click **OK** to add the rule to the table.

Step 4 Click **Apply** to save the access rule to your configuration.

Access Rule Properties

When you add or edit an access rule, you can configure the following properties. In many fields, you can click the "..." button on the right of the edit box to select, create, or edit objects that are available for the field.

Interface

The interface to which the rule applies. Select Any to create a global rule. For bridge groups in routed mode, you can select either the Bridge Virtual Interface (BVI) or the bridge group member interfaces.

Action: Permit/Deny

Whether you are permitting (allowing) the described traffic or are denying (dropping) it.

Source Criteria

The characteristics of the originator of the traffic you are trying to match. You must configure Source, but the other properties are optional.

Source

The IPv4 or IPv6 address of the source. The default is **any**, which matches all IPv4 or IPv6 addresses; you can use **any4** to target IPv4 only, or **any6** to target IPv6 only. You can specify a single host address (such as 10.100.10.5 or 2001:DB8::0DB8:800:200C:417A), a subnet (in 10.100.10.0/24 or 10.100.10.0/255.255.255.0 format, or for IPv6, 2001:DB8:0:CD30::/60), the name of a network object or network object group, or the name of an interface.

User

If you enable the identity firewall, you can specify a user or user group as the traffic source. The IP address the user is currently using will match the rule. You can specify a username (DOMAIN\user), a user group (DOMAIN\\group, note the double \ indicates a group name), or a user object group. For this field, it is far easier to click "**...**" to select names from your AAA server group than to type them in.

Security Group

If you enable Cisco Trustsec, you can specify a security group name or tag (1-65533), or security group object.

More Options > Source Service

If you specify TCP, UDP, or SCTP as the destination service, you can optionally specify a predefined service object for TCP, UDP, TCP-UDP, or SCTP, or use your own object. Typically, you define the destination service only and not the source service. Note that if you define the source service, the destination service protocol must match it (for example, both TCP, with or without port definitions).

Destination Criteria

The characteristics of the target of the traffic you are trying to match. You must configure Destination, but the other properties are optional.

Destination

The IPv4 or IPv6 address of the destination. The default is **any**, which matches all IPv4 or IPv6 addresses; you can use **any4** to target IPv4 only, or **any6** to target IPv6 only. You can specify a

single host address (such as 10.100.10.5 or 2001:DB8::0DB8:800:200C:417A), a subnet (in 10.100.10.0/24 or 10.100.10.0/255.255.255.0 format, or for IPv6, 2001:DB8:0:CD30::/60), the name of a network object or network object group, or the name of an interface.

Security Group

If you enable Cisco Trustsec, you can specify a security group name or tag (1-65533), or security group object.

Service

The protocol of the traffic, such as IP, TCP, UDP, and optionally ports for TCP, UDP, or SCTP. The default is IP, but you can select a more specific protocol to target traffic with more granularity. Typically, you would select some type of service object. For TCP, UDP, and SCTP, you can specify ports, for example, tcp/80, tcp/http, tcp/10-20 (for a range of ports), tcp-udp/80 (match any TCP or UDP traffic on port 80), sctp/diameter, and so forth.

Description

A explanation of the purpose of the rule, up to 100 characters per line. You can enter multiple lines; each line is added as a remark in the CLI, and the remarks are placed before the rule.

Note If you add remarks with non-English characters on one platform (such as Windows) then try to remove them from another platform (such as Linux), you might not be able to edit or delete them because the original characters might not be correctly recognized. This limitation is due to an underlying platform dependency that encodes different language characters in different ways.

Enable Logging; Logging Level; More Options > Logging Interval

The logging options define how syslog messages will be generated for rules. You can implement the following logging options:

Deselect Enable Logging

This will disable logging for the rule. No syslog messages of any type will be issued for connections that match this rule.

Select Enable Logging with Logging Level = Default

This provides the default logging for rules. Syslog message 106023 is issued for each denied connection. If the appliance comes under attack, the frequency of issuing this message could impact services.

Select Enable Logging with Non-Default Logging Level

This provides a summarized syslog message, 106100, instead of 106023. Message 106100 is issued upon first hit, then again at each interval configured in **More Options > Logging Interval** (default is every 300 seconds, you can specify 1-600), showing the number of hits during that interval. The recommended logging level is **Informational**.

Summarizing deny messages can reduce the impact of attacks and possibly make it easier for you to analyze messages. If you do come under a denial of service attack, you might see message 106101, which indicates that the number of cached deny flows used to produce the hit count for message 106100 has exceeded the maximum for an interval. At this point, the appliance stops collecting statistics until the next interval to mitigate the attack.

More Options > Traffic Direction

Whether the rule is for the **In** or **Out** direction. **In** is the default, and it is the only option for global and management access rules.

More Options > Enable Rule

Whether the rule is active on the device. Disabled rules appear with strike-through text in the rule table. Disabling a rule lets you stop its application to traffic without deleting it, so you can enable it again later if you decide you need it.

More Options > Time Range

The name of the time range object that defines the times of day and days of the week when the rule should be active. If you do not specify a time range, the rule is always active.

Configure Advanced Options for Access Rules

Advanced access rule options allow you to customize certain aspects of rule behavior, but these options have defaults that are appropriate in most cases.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Access Rules**.

Step 2 Click the **Advanced** button below the rule table.

Step 3 Configure the following options as required:

- **Advanced Logging Settings**—If you configure non-default logging, the system caches deny flows to develop statistics for message 106100, as explained in Evaluating Syslog Messages for Access Rules, on page 27. To prevent unlimited consumption of memory and CPU resources, the ASA places a limit on the number of concurrent *deny* flows because they can indicate an attack. Message 106101 is issued when the limit is reached. You can control the following aspects related to 106101.
  - Maximum Deny-flows—The maximum number of deny flows permitted before the ASA stops caching flows, between 1 and 4096. The default is 4096.
  - Alert Interval—The amount of time (1-3600 seconds) between issuing system log message 106101, which indicates that the maximum number of deny flows was reached. The default is 300 seconds.
- **Per User Override table**—Whether to allow a dynamic user ACL that is downloaded for user authorization from a RADIUS server to override the ACL assigned to the interface. For example, if the interface ACL denies all traffic from 10.0.0.0, but the dynamic ACL permits all traffic from 10.0.0.0, then the dynamic ACL overrides the interface ACL for that user. Check the **Per User Override** box for each interface that should allow user overrides (inbound direction only). If the per user override feature is disabled, the access rule provided by the RADIUS server is combined with the access rule configured on that interface.

  By default, VPN remote access traffic is not matched against interface ACLs. However, if you deselect the **Enable inbound VPN sessions to bypass interface access lists** setting on the Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Client Connection Profiles pane), the behavior depends on whether there is a VPN filter applied in the group policy (see the Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add/Edit > General > More Options > Filter field) and whether you set the Per User Override option:

- No Per User Override, no VPN filter —Traffic is matched against the interface ACL.
- No Per User Override, VPN filter —Traffic is matched first against the interface ACL, then against the VPN filter.
- Per User Override, VPN filter —Traffic is matched against the VPN filter only.

- **Object Group Search Setting**—You can reduce the memory required to search access rules that use object groups by selecting **Enable Object Group Search Algorithm**, but this is at the expense of rule lookup performance. When enabled, object group search does not expand network objects, but instead searches access rules for matches based on those group definitions.

  Select **Enable Object Group Search Threshold** to set a threshold to help prevent performance degradation. When operating with a threshold, for each connection, both the source and destination IP addresses are matched against network objects. If the number of objects matched by the source address times the number matched by the destination address exceeds 10,000, the connection is dropped. Configure your rules to prevent an excessive number of matches.

  Note
  Object group search works with network and service objects only. It does not work with security group objects. Do not enable this feature if the ACLs include security groups. The result can be inactive ACLs or other unexpected behavior.

- **Forward Reference Setting**—Normally, you cannot reference an object or object group that does not exist in an ACL or object group, or delete one that is currently referenced. You also cannot reference an ACL that does not exist in an **access-group** command (to apply access rules). However, you can change this default behavior so that you can "forward reference" objects or ACLs before you create them. Until you create the objects or ACLs, any rules or access groups that reference them are ignored. Select **Enable the forward reference of objects and object-groups** to enable forward referencing. Be aware that if you enable forward referencing, ASDM cannot tell the difference between a typo reference to an existing object and a forward reference.

Step 4 Click **OK**.

Configure Management Access Rules

You can configure an interface ACL that controls to-the-box management traffic from a specific peer (or set of peers) to the ASA. One scenario in which this type of ACL would be useful is when you want to block IKE Denial of Service attacks.

Unlike regular access rules, there is no implicit deny at the end of a set of management rules for an interface. Instead, any connection that does not match a management access rule is then evaluated by regular access control rules.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **Management Access Rules**.

The rules are organized by interface. Each group is equivalent to the extended ACL that is created and assigned to the interface as a control plane ACL. These ACLs also appear on the Access Rules and ACL Manager pages.

Step 2 Do any of the following:

- To add a new rule, choose **Add > Add Management Access Rule**.
- To insert a rule at a specific location within a container, select an existing rule and choose **Add > Insert** to add the rule above it, or choose **Add > Insert After**.
- To edit a rule, select it and click **Edit**.

Step 3 Fill in the rule properties. The primary options to select are:

- **Interface**—The interface to which the rule applies. For bridge groups in routed mode, you can create access rules for both the Bridge Virtual Interface (BVI) and each bridge group member interface.
- **Action: Permit/Deny**—Whether you are permitting (allowing) the described traffic or are denying (dropping) it.
- **Source/Destination criteria**—A definition of the source (originating address) and destination (target address of the traffic flow). You typically configure IPv4 or IPv6 addresses of hosts or subnets, which you can represent with network or network object groups. You can also specify a user or user group name for the source. Additionally, you can use the Service field to identify the specific type of traffic if you want to focus the rule more narrowly than all IP traffic. If you implement Trustsec, you can use security groups to define source and destination.

For detailed information on all of the available options, see Access Rule Properties, on page 20.

When you are finished defining the rule, click **OK** to add the rule to the table.

Step 4 Click **Apply** to save the rule to your configuration.

Configure EtherType Rules

EtherType rules apply to non-IP layer-2 traffic on bridge group member interfaces (in routed or transparent firewall mode). You can use these rules to permit or drop traffic based on the EtherType value in the layer-2 packet. With EtherType rules, you can control the flow of non-IP traffic across the ASA.

You can apply both extended and EtherType access rules to a bridge group member interface. EtherType rules take precedence over the extended access rules.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **EtherType Rules**.

The rules are organized by interface and direction. Each group is equivalent to the EtherType ACL that is created and assigned to the interface.

Step 2 Do any of the following:

- To add a new rule, choose **Add** > **Add EtherType Rule**.
- To insert a rule at a specific location within a container, select an existing rule and choose **Add** > **Insert** to add the rule above it, or choose **Add** > **Insert After**.
- To edit a rule, select it and click **Edit**.

Step 3 Fill in the rule properties. The primary options to select are:

- **Interface**—The interface to which the rule applies.
- **Action: Permit/Deny**—Whether you are permitting (allowing) the described traffic or are denying (dropping) it.
- EtherType—You can match traffic using the following options:
  - **any**—Matches all traffic.
  - **bpdu**—Bridge protocol data units, which are allowed by default. When you apply the configuration, this keyword is converted to **dsap bpdu** on the device.
  - **dsap**—The IEEE 802.2 Logical Link Control packet's Destination Service Access Point address. You also need to include the address you want to permit or deny in hexadecimal, from 0x01 to 0xff, in **DSAP Value**. Following are the values for some common addresses:
    - **0x42**—bridge protocol data units (BPDU). When you apply the configuration, this is converted to **dsap bpdu** on the device.
    - **0xe0**—Internet Packet Exchange (IPX) 802.2 LLC. When you apply the configuration, this is converted to **dsap ipx** on the device.
    - **0xfe**—Intermediate System to Intermediate System (IS-IS). When you apply the configuration, this is converted to **dsap isis** on the device.
    - **0xff**—raw IPX 802.3 format. When you apply the configuration, this is converted to **dsap raw-ipx** on the device.
  - **eii-ipx**—Ethernet II IPX format, EtherType 0x8137.
  - **ipx**—Internet Packet Exchange (IPX). This keyword is a shortcut for configuring three separate rules, for **dsap ipx**, **dsap raw-ipx**, and **eii-ipx**. The conversion is made when you apply the configuration to the device.
  - **isis**—Intermediate System to Intermediate System (IS-IS). When you apply the configuration, this keyword is converted to **dsap isis** on the device.
  - **mpls-multicast**— MPLS multicast.
  - **mpls-unicast**—MPLS unicast.
  - *hex_number*—Any EtherType that can be identified by a 16-bit hexadecimal number 0x600 to 0xffff. See RFC 1700, “Assigned Numbers,” at http://www.ietf.org/rfc/rfc1700.txt for a list of EtherTypes.
- **Description**—A explanation of the purpose of the rule, up to 100 characters per line. You can enter multiple lines; each line is added as a remark in the CLI, and the remarks are placed before the rule.
- **More Options** > **Direction**—Whether the rule is for the **In** or **Out** direction. **In** is the default.

When you are finished defining the rule, click **OK** to add the rule to the table.

Step 4 Click **Apply** to save the rule to your configuration.

Configure ICMP Access Rules

By default, you can send ICMP packets to any interface using either IPv4 or IPv6, with these exceptions:

- The ASA does not respond to ICMP echo requests directed to a broadcast address.
- The ASA only responds to ICMP traffic sent to the interface that traffic comes in on; you cannot send ICMP traffic through an interface to a far interface.

To protect the device from attacks, you can use ICMP rules to limit ICMP access to interfaces to particular hosts, networks, or ICMP types. ICMP rules function like access rules, where the rules are ordered, and the first rule that matches a packet defines the action.

If you configure any ICMP rule for an interface, an implicit deny ICMP rule is added to the end of the ICMP rule list, changing the default behavior. Thus, if you want to simply deny a few message types, you must include a permit any rule at the end of the ICMP rule list to allow the remaining message types.

We recommend that you always grant permission for the ICMP unreachable message type (type 3). Denying ICMP unreachable messages disables ICMP path MTU discovery, which can halt IPsec and PPTP traffic. Additionally ICMP packets in IPv6 are used in the IPv6 neighbor discovery process.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **ICMP**.

Step 2 Configure ICMP rules:

a) Add a rule (**Add > Rule**, **Add > IPv6 Rule**, or **Add > Insert**), or select a rule and edit it.
b) Select the ICMP type you want to control, or **any** to apply to all types.
c) Select the interface to which the rule applies. You must create separate rules for each interface.
d) Select whether you are permitting or denying access for matching traffic.
e) Select **Any Address** to apply the rule to all traffic. Alternatively, enter the address and mask (for IPv4) or address and prefix length (for IPv6) of the host or network you are trying to control.
f) Click **OK**.

Step 3 (Optional) To set ICMP unreachable message limits, set the following options. Increasing the rate limit, along with enabling the **Decrement time to live for a connection** option in a service policy (on the Configuration > Firewall > Service Policy Rules > Rule Actions > Connection Settings dialog box), is required to allow a trace route through the ASA that shows the ASA as one of the hops.

- Rate Limit—Sets the rate limit of unreachable messages, between 1 and 100 messages per second. The default is 1 message per second.
- Burst Size—Sets the burst rate, between 1 and 10. The system sends this number of replies, but subsequent replies are not sent until the rate limit is reached.

Step 4 Click **Apply**.

Monitoring Access Rules

The Access Rules page includes hit counts for each rule. Mouse over the hit count to see the update time and interval for the count. To reset the hit count, right click the rule and select **Clear Hit Count**, but be aware that this clears the count for all rules applied to the same interface in the same direction.

Evaluating Syslog Messages for Access Rules

Use a syslog event viewer, such as the one in ASDM, to view messages related to access rules.

If you use default logging, you see syslog message 106023 for explicitly denied flows only. Traffic that matches the "implicit deny" entry that ends the rule list is not logged.

If the ASA is attacked, the number of syslog messages for denied packets can be very large. We recommend that you instead enable logging using syslog message 106100, which provides statistics for each rule (including permit rules) and enables you to limit the number of syslog messages produced. Alternatively, you can disable all logging for a given rule.

When you enable logging for message 106100, if a packet matches an ACE, the ASA creates a flow entry to track the number of packets received within a specific interval. The ASA generates a syslog message at the first hit and at the end of each interval, identifying the total number of hits during the interval and the time stamp for the last hit. At the end of each interval, the ASA resets the hit count to 0. If no packets match the ACE during an interval, the ASA deletes the flow entry. When you configure logging for a rule, you can control the interval and even the severity level of the log message, per rule.

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.

Permitted packets that belong to established connections do not need to be checked against ACLs; only the initial packet is logged and included in the hit count. For connectionless protocols, such as ICMP, all packets are logged, even if they are permitted, and all denied packets are logged.

See the *syslog messages guide* for detailed information about these messages.

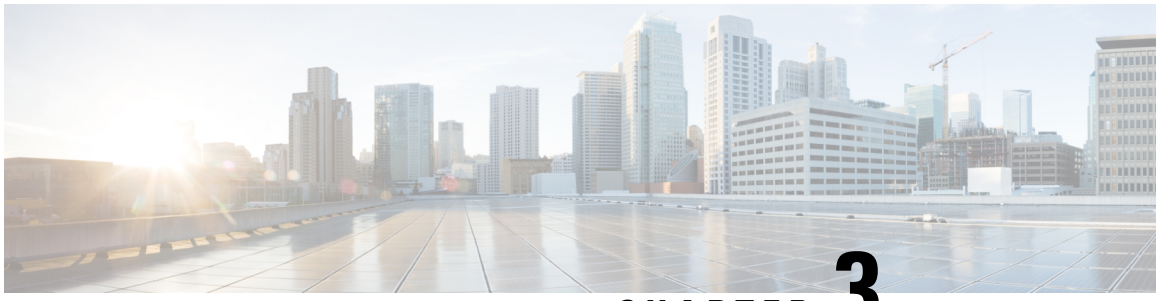


Tip When you enable logging for message 106100, if a packet matches an ACE, the ASA creates a flow entry to track the number of packets received within a specific interval. The ASA has a maximum of 32 K logging flows for ACEs. A large number of flows can exist concurrently at any point of time. To prevent unlimited consumption of memory and CPU resources, the ASA places a limit on the number of concurrent *deny* flows; the limit is placed on deny flows only (not on permit flows) because they can indicate an attack. When the limit is reached, the ASA does not create a new deny flow for logging until the existing flows expire, and issues message 106101. You can control the frequency of this message, and the maximum number of deny flows cached, in the advanced settings; see Configure Advanced Options for Access Rules, on page 22.

History for Access Rules

Feature Name	Platform Releases	Description
Interface access rules	7.0(1)	Controlling network access through the ASA using ACLs. We introduced the following screen: Configuration > Firewall > Access Rules.
Global access rules	8.3(1)	Global access rules were introduced. We modified the following screen: Configuration > Firewall > Access Rules.
Support for Identity Firewall	8.4(2)	You can now use identity firewall users and groups for the source and destination. You can use an identity firewall ACL with access rules, AAA rules, and for VPN authentication.
EtherType ACL support for IS-IS traffic	8.4(5), 9.1(2)	In transparent firewall mode, the ASA can now pass IS-IS traffic using an EtherType ACL. We modified the following screen: Configuration > Device Management > Management Access > EtherType Rules.
Support for TrustSec	9.0(1)	You can now use TrustSec security groups for the source and destination. You can use an identity firewall ACL with access rules.
Unified ACL for IPv4 and IPv6	9.0(1)	ACLs now support IPv4 and IPv6 addresses. You can even specify a mix of IPv4 and IPv6 addresses for the source and destination. The **any** keyword was changed to represent IPv4 and IPv6 traffic. The **any4** and **any6** keywords were added to represent IPv4-only and IPv6-only traffic, respectively. The IPv6-specific ACLs are deprecated. Existing IPv6 ACLs are migrated to extended ACLs. See the release notes for more information about migration. We modified the following screens: Configuration > Firewall > Access Rules Configuration > Remote Access VPN > Network (Client) Access > Group Policies > General > More Options
Extended ACL and object enhancement to filter ICMP traffic by ICMP code	9.0(1)	ICMP traffic can now be permitted/denied based on ICMP code. We introduced or modified the following screens: Configuration > Firewall > Objects > Service Objects/Groups Configuration > Firewall > Access Rule
Transactional Commit Model on Access Group Rule Engine	9.1(5)	When enabled, a rule update is applied after the rule compilation is completed; without affecting the rule matching performance. We introduced the following screen: Configuration > Device Management > Advanced > Rule Engine.

Feature Name	Platform Releases	Description
Configuration session for editing ACLs and objects. Forward referencing of objects and ACLs in access rules.	9.3(2)	You can now edit ACLs and objects in an isolated configuration session. You can also forward reference objects and ACLs, that is, configure rules and access groups for objects or ACLs that do not yet exist.
Access rule support for Stream Control Transmission Protocol (SCTP)	9.5(2)	You can now create access rules using the **sctp** protocol, including port specifications. We modified the add/edit dialog boxes for access rules on the **Configuration** > **Firewall** > **Access Rules** page.
Ethertype rule support for the IEEE 802.2 Logical Link Control packet's Destination Service Access Point address.	9.6(2)	You can now write Ethertype access control rules for the IEEE 802.2 Logical Link Control packet's Destination Service Access Point address. Because of this addition, the **bpdu** keyword no longer matches the intended traffic. Rewrite **bpdu** rules for **dsap 0x42**. We modified the following screen: **Configuration** > **Firewall** > **EtherType Rules**.
Support in routed mode for Ethertype rules on bridge group member interfaces and extended access rules on Bridge Group Virtual Interfaces (BVI).	9.7(1)	You can now create Ethertype ACLs and apply them to bridge group member interfaces in routed mode. You can also apply extended access rules to the Bridge Virtual Interface (BVI) in addition to the member interfaces. We modified the following screens: **Configuration** > **Firewall** > **Access Rules**, **Configuration** > **Firewall** > **Ethertype Rules**.
EtherType access control list changes.	9.9(1)	EtherType access control lists now support Ethernet II IPX (EII IPX). In addition, new keywords are added to the DSAP keyword to support common DSAP values: BPDU (0x42), IPX (0xE0), Raw IPX (0xFF), and ISIS (0xFE). Consequently, existing EtherType access control entries that use the BPDU or ISIS keywords will be converted automatically to use the DSAP specification, and rules for IPX will be converted to 3 rules (DSAP IPX, DSAP Raw IPX, and EII IPX). In addition, packet capture that uses IPX as an EtherType value has been deprecated, because IPX corresponds to 3 separate EtherTypes. We modified the following screens: **Configuration** > **Firewall** > **Ethertype Rules**.
The object group search threshold is now disabled by default.	9.12(1)	If you enabled object group search, the feature was subject to a threshold to help prevent performance degradation. That threshold is now disabled by default. You can enable it by using the **object-group-search threshold** command. We changed the following screen: **Configuration** > **Access Rules** > **Advanced**.



CHAPTER 3

Objects for Access Control

Objects are reusable components for use in your configuration. You can define and use them in ASA configurations in the place of inline IP addresses, services, names, and so on. Objects make it easy to maintain your configurations because you can modify an object in one place and have it be reflected in all other places that are referencing it. Without objects you would have to modify the parameters for every feature when required, instead of just once. For example, if a network object defines an IP address and subnet mask, and you want to change the address, you only need to change it in the object definition, not in every feature that refers to that IP address.

- Guidelines for Objects, on page 31
- Configure Objects, on page 32
- Monitoring Objects, on page 38
- History for Objects, on page 38

Guidelines for Objects

IPv6 Guidelines

Supports IPv6 with the following restrictions:

- You can mix IPv4 and IPv6 entries in a network object group, but you cannot use a mixed object group for NAT.

Additional Guidelines and Limitations

- Objects must have unique names, because objects and object groups share the same name space. While you might want to create a network object group named “Engineering” and a service object group named “Engineering,” you need to add an identifier (or “tag”) to the end of at least one object group name to make it unique. For example, you can use the names “Engineering_admins” and “Engineering_hosts” to make the object group names unique and to aid in identification.
- If you enter more than one item in source or destination address, or source or destination service, in an ACL or access rule, ASDM automatically creates an object group for them with the prefix DM_INLINE. These objects are not shown on the objects page, but they are defined on the device.
- Object names are limited to 64 characters, including letters, numbers, and these characters: .!@#$%^&()-_{}. Object names are case-sensitive.

- You cannot remove an object or make an object empty if it is used in a command, unless you enable forward referencing (in the access rules advanced settings).

Configure Objects

The following sections describe how to configure objects that are primarily used on access control.

Configure Network Objects and Groups

Network objects and groups identify IP addresses or host names. Use these objects in access control lists to simplify your rules.

Configure a Network Object

A network object can contain a host, a network IP address, a range of IP addresses, or a fully qualified domain name (FQDN).

You can also enable NAT rules on the object (excepting FQDN objects). For more information about configuring object NAT, see Network Address Translation (NAT), on page 149.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Network Objects/Group**.

Step 2 Do one of the following:

- Choose **Add > Network Object** to add a new object. Enter a name and optionally, a description.
- Choose an existing object and click **Edit**.

Step 3 Configure the address for the object based on the object **Type** and **IP version** fields.

- **Host**—The IPv4 or IPv6 address of a single host. For example, 10.1.1.1 or 2001:DB8::0DB8:800:200C:417A.
- **Network**—The address of a network. For IPv4, include the mask, for example, **IP address** = 10.0.0.0 **Netmask** = 255.0.0.0. For IPv6, include the prefix, such as **IP Address** = 2001:DB8:0:CD30:: **Prefix Length** = 60.
- **Range**—A range of addresses. You can specify IPv4 or IPv6 ranges. Do not include masks or prefixes.
- **FQDN**—A fully-qualified domain name, that is, the name of a host, such as www.example.com.

Step 4 Click **OK**, then click **Apply**.

You can now use this network object when you create a rule. If you edit an object, the change is inherited automatically by any rules using the object.

Configure a Network Object Group

Network object groups can contain multiple network objects as well as inline networks or hosts. Network object groups can include a mix of both IPv4 and IPv6 addresses.

However, you cannot use a mixed IPv4 and IPv6 object group for NAT, or object groups that include FQDN objects.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups**.

Step 2 Do one of the following:

- Choose **Add > Network Object Group** to add a new object. Enter a name and optionally, a description.
- Choose an existing object and click **Edit**.

Step 3 Add network objects to the group using any combination of the following techniques:

- **Existing Network Objects/Groups**—Select any already defined network object or group and click **Add** to include them in the group.
- **Create New Network Object Member**—Enter the criteria for a new network object and click **Add**.

Step 4 After you add all the member objects, click **OK**, then click **Apply**.

You can now use this network object group when you create a rule. For an edited object group, the change is inherited automatically by any rules using the group.

Configure Service Objects and Service Groups

Service objects and groups identify protocols and ports. Use these objects in access control lists to simplify your rules.

Configure a Service Object

A service object can contain a single protocol specification.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Service Object/Group**.

Step 2 Do one of the following:

- Choose **Add > Service Object** to add a new object. Enter a name and optionally, a description.
- Choose an existing object and click **Edit**.

Step 3 Choose the service type and fill in details as needed:

- Protocol—A number between 0-255, or a well-known name, such as **ip**, **tcp**, **udp**, **gre**, and so forth..
- ICMP, ICMP6—You can leave the message type and code fields blank to match any ICMP/ICMP version 6 message. You can optionally specify the ICMP type by name or number (0-255) to limit the object to that message type. If you specify a type, you can optionally specify an ICMP code for that type (1-255). If you do not specify the code, then all codes are used.
- TCP, UDP, SCTP—You can optionally specify ports for the source, destination, or both. You can specify the port by name or number. You can include the following operators:
  - **<**—Less than. For example, <80.
  - **>**—Greater than. For example, >80.
  - **!=**—Not equal to. For example, !=80.
  - **-** (hyphen)—An inclusive range of values. For example, 100-200.

Step 4 Click **OK**, and then **Apply**.

Configure a Service Group

A service object group includes a mix of protocols, if desired, including optional source and destination ports for protocols that use them, and ICMP type and code.

Before you begin

You can model all services using the generic service object group, which is explained here. However, you can still configure the types of service group objects that were available prior to ASA 8.3(1). These legacy objects include TCP/UDP/TCP-UDP port groups, protocol groups, and ICMP groups. The contents of these groups are equivalent to the associated configuration in the generic service object group, with the exception of ICMP groups, which do not support ICMP6 or ICMP codes. If you still want to use these legacy objects, for detailed instructions, see the **object-service** command description in the command reference on Cisco.com.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Service Objects/Groups**.

Step 2 Do one of the following:

- Choose **Add > Service Group** to add a new object. Enter a name and optionally, a description.
- Choose an existing object and click **Edit**.

Step 3 Add service objects to the group using any combination of the following techniques:

- **Existing Service/Service Group**—Select any already defined service, service object, or group and click **Add** to include them in the group.
- **Create New Member**—Enter the criteria for a new service object and click **Add**. If you give the object a name, when you apply changes, the new object is created and added to the group; otherwise, unnamed

objects are members of this group only. You cannot name TCP-UDP objects; these are members of the group only.

Step 4 After you add all the member objects, click **OK**, then click **Apply**.

You can now use this service object group when you create a rule. For an edited object group, the change is inherited automatically by any rules using the group.

Configure Local User Groups

You can create local user groups for use in features that support the identity firewall by including the group in an extended ACL, which in turn can be used in an access rule, for example.

The ASA sends an LDAP query to the Active Directory server for user groups globally defined in the Active Directory domain controller. The ASA imports these groups for identity-based rules. However, the ASA might have localized network resources that are not defined globally that require local user groups with localized security policies. Local user groups can contain nested groups and user groups that are imported from Active Directory. The ASA consolidates local and Active Directory groups.

A user can belong to local user groups and user groups imported from Active Directory.

Because you can use usernames and user group names directly in an ACL, you need to configure local user groups only if:

- You want to create a group of users defined in the LOCAL database.
- You want to create a group of users or user groups that are not captured in a single user group defined on the AD server.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Local User Groups**.

Step 2 Do one of the following:

- Choose **Add** to add a new object. Enter a name and optionally, a description.
- Choose an existing object and click **Edit**.

Step 3 Add users or groups to the object using any of these methods:

- **Select existing users or groups**—Select the domain that contains the user or group, then pick the user or group name from the lists and click **Add**. For long lists, use the Find box to help locate the user. The names are pulled from the server for the selected domain.
- **Manually type user names**—You can simply type in the user or group names in the bottom edit box and click **Add**. When using this method, the selected domain name is ignored, and the default domain is used if you do not specify one. For users, the format is *domain_name**username*; for groups, there is a double \\, *domain_name**group_name*.

Step 4 After you add all the member objects, click **OK**, then click **Apply**.

You can now use this user object group when you create a rule. For an edited object group, the change is inherited automatically by any rules using the group.

Configure Security Group Object Groups

You can create security group object groups for use in features that support Cisco TrustSec by including the group in an extended ACL, which in turn can be used in an access rule, for example.

When integrated with Cisco TrustSec, the ASA downloads security group information from the ISE. The ISE acts as an identity repository, by providing Cisco TrustSec tag-to-user identity mapping and Cisco TrustSec tag-to-server resource mapping. You provision and manage security group ACLs centrally on the ISE.

However, the ASA might have localized network resources that are not defined globally that require local security groups with localized security policies. Local security groups can contain nested security groups that are downloaded from the ISE. The ASA consolidates local and central security groups.

To create local security groups on the ASA, you create a local security object group. A local security object group can contain one or more nested security object groups or Security IDs or security group names. You can also create a new Security ID or security group name that does not exist on the ASA.

You can use the security object groups you create on the ASA to control access to network resources. You can use the security object group as part of an access group or service policy.



Tip If you create a group with tags or names that are not known to the ASA, any rules that use the group will be inactive until the tags or names are resolved with ISE.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Security Group Object Groups**.

Step 2 Do one of the following:

- Choose **Add** to add a new object. Enter a name and optionally, a description.
- Choose an existing object and click **Edit**.

Step 3 Add security groups to the object using any of these methods:

- **Select existing local security group object groups**—Pick from the list of objects already defined and click **Add**. For long lists, use the Find box to help locate the object.
- **Select security groups discovered from ISE**—Pick groups from the list of existing groups and click **Add**.
- **Manually add security tags or names**—You can simply type in the tag number or security group name in the bottom edit box and click **Add**. A tag is a number from 1 to 65533 and is assigned to a device through IEEE 802.1X authentication, web authentication, or MAC authentication bypass (MAB) by the ISE. Security group names are created on the ISE and provide user-friendly names for security groups.

The security group table maps SGTs to security group names. Consult your ISE configuration for the valid tags and names.

Step 4 After you add all the member objects, click **OK**, then click **Apply**.

You can now use this security group object group when you create a rule. For an edited object group, the change is inherited automatically by any rules using the group.

Configure Time Ranges

A time range object defines a specific time consisting of a start time, an end time, and optional recurring entries. You use these objects on ACL rules to provide time-based access to certain features or assets. For example, you could create an access rule that allows access to a particular server during working hours only.



Note You can include multiple periodic entries in a time range object. If a time range has both absolute and periodic values specified, then the periodic values are evaluated only after the absolute start time is reached, and they are not further evaluated after the absolute end time is reached.

Creating a time range does not restrict access to the device. This procedure defines the time range only. You must then use the object in an access control rule.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Time Ranges**.

Step 2 Do one of the following:

- Choose **Add** to add a new time range. Enter a name and optionally, a description.
- Choose an existing time range and click **Edit**.

Step 3 Choose the overall start and end time.

The default is to start now and never end, but you can set specific dates and times. The time range is inclusive of the times that you enter.

Step 4 (Optional) Configure recurring periods within the overall active time, such as the days of the week or the recurring weekly interval in which the time range will be active.

a) Click **Add**, or select an existing period and click **Edit**.

b) Do one of the following:

- Click **Specify days of the week and times on which this recurring range will be active**, and choose the days and times from the lists.
- Click **Specify a weekly interval when this recurring range will be active**, and choose the days and times from the lists.

c) Click **OK**.

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

Monitoring Objects

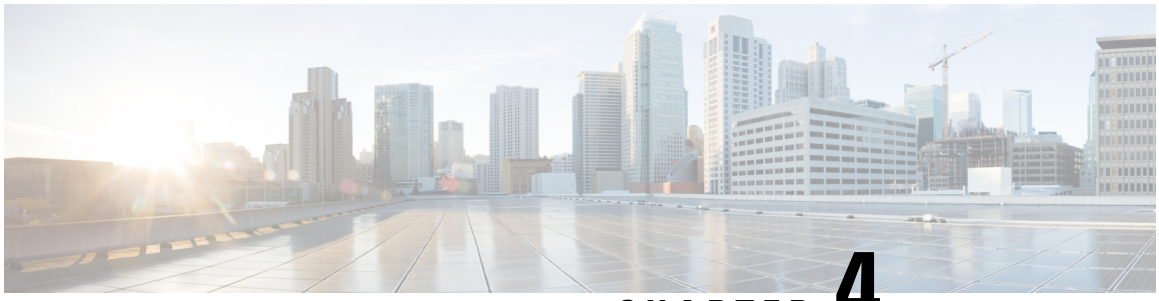
For network, service, and security group objects, you can analyze the usage of an individual object. From their page in the **Configuration > Firewall > Objects** folder, click the **Where Used** button.

For network objects, you can also click the Not Used button to find objects that are not used in any rules or other objects. This display gives you a short-cut for deleting these unused objects.

History for Objects

Feature Name	Platform Releases	Description
Object groups	7.0(1)	Object groups simplify ACL creation and maintenance.
Regular expressions and policy maps	7.2(1)	Regular expressions and policy maps were introduced to be used under inspection policy maps. The following commands were introduced: **class-map type regex**, **regex**, **match regex**.
Objects	8.3(1)	Object support was introduced.
User Object Groups for Identity Firewall	8.4(2)	User object groups for identity firewall were introduced.
Security Group Object Groups for Cisco TrustSec	8.4(2)	Security group object groups for Cisco TrustSec were introduced.
Mixed IPv4 and IPv6 network object groups	9.0(1)	Previously, network object groups could only contain all IPv4 addresses or all IPv6 addresses. Now network object groups can support a mix of both IPv4 and IPv6 addresses. **Note** You cannot use a mixed object group for NAT.
Extended ACL and object enhancement to filter ICMP traffic by ICMP code	9.0(1)	ICMP traffic can now be permitted/denied based on ICMP code. We introduced or modified the following screens: Configuration > Firewall > Objects > Service Objects/Groups, Configuration > Firewall > Access Rule

Feature Name	Platform Releases	Description
Service object support for Stream Control Transmission Protocol (SCTP)	9.5(2)	You can now create service objects and groups that specific SCTP ports. We modified the add/edit dialog boxes for service objects and groups on the **Configuration** > **Firewall** > **Objects** > **Service Objects/Groups** page.



CHAPTER 4

Access Control Lists

Access control lists (ACLs) are used by many different features. When applied to interfaces or globally as access rules, they permit or deny traffic that flows through the appliance. For other features, the ACL selects the traffic to which the feature will apply, performing a matching service rather than a control service.

The following sections explain the basics of ACLs and how to configure and monitor them. Access rules, ACLs applied globally or to interfaces, are explained in more detail in Access Rules, on page 11.

- About ACLs, on page 41
- Licensing for Access Control Lists, on page 45
- Guidelines for ACLs, on page 46
- Configure ACLs, on page 47
- Monitoring ACLs, on page 54
- History for ACLs, on page 55

About ACLs

Access control lists (ACLs) identify traffic flows by one or more characteristics, including source and destination IP address, IP protocol, ports, EtherType, and other parameters, depending on the type of ACL. ACLs are used in a variety of features. ACLs are made up of one or more access control entries (ACEs).

ACL Types

The ASA uses the following types of ACLs:

- Extended ACLs—Extended ACLs are the main type that you will use. These ACLs are used for access rules to permit and deny traffic through the device, and for traffic matching by many features, including service policies, AAA rules, WCCP, Botnet Traffic Filter, and VPN group and DAP policies. In ASDM, many of these features have their own rules pages and they cannot use extended ACLs that you define in the ACL Manager, although ACL Manager will display the ACLs created on those pages. See Configure Extended ACLs, on page 47.
- EtherType ACLs—EtherType ACLs apply to non-IP layer-2 traffic on bridge group member interfaces only. You can use these rules to permit or drop traffic based on the EtherType value in the layer-2 packet. With EtherType ACLs, you can control the flow of non-IP traffic across the device. See Configure EtherType Rules, on page 24.

- Webtype ACLs—Webtype ACLs are used for filtering clientless SSL VPN traffic. These ACLs can deny access based on URLs or destination addresses. See Configure Webtype ACLs, on page 52.
- Standard ACLs—Standard ACLs identify traffic by destination address only. There are few features that use them: route maps and VPN filters. Because VPN filters also allow extended access lists, limit standard ACL use to route maps. See Configure Standard ACLs, on page 51.

The following table lists some common uses for ACLs and the type to use.

Table 2: ACL Types and Common Uses

ACL Use	ACL Type	Description
Control network access for IP traffic (routed and transparent mode)	Extended	The ASA does not allow any traffic from a lower security interface to a higher security interface unless it is explicitly permitted by an extended ACL. In routed mode, you must use an ACL to permit traffic between a bridge group member interface and an interface outside same the bridge group. **Note** To access the ASA interface for management access, you do not also need an ACL allowing the host IP address. You only need to configure management access according to the general operations configuration guide.
Identify traffic for AAA rules	Extended	AAA rules use ACLs to identify traffic.
Augment network access control for IP traffic for a given user	Extended, downloaded from a AAA server per user	You can configure the RADIUS server to download a dynamic ACL to be applied to the user, or the server can send the name of an ACL that you already configured on the ASA.
VPN access and filtering	Extended Standard	Group policies for remote access and site to site VPNs use standard or extended ACLs for filtering. Remote access VPNs also use extended ACLs for client firewall configurations and dynamic access policies.
Identify traffic in a traffic class map for Modular Policy Framework	Extended	ACLs can be used to identify traffic in a class map, which is used for features that support Modular Policy Framework. Features that support Modular Policy Framework include TCP and general connection settings, and inspection.
For bridge group member interfaces, control network access for non-IP traffic	EtherType	You can configure an ACL that controls traffic based on its EtherType for any interface that is a member of a bridge group.
Identify route filtering and redistribution	Standard Extended	Various routing protocols use standard ACLs for route filtering and redistribution (through route maps) for IPv4 addresses, and extended ACLs for IPv6.
Filtering for clientless SSL VPN	Webtype	You can configure a webtype ACL to filter URLs and destinations.

The ACL Manager

The ACL Manager appears in two forms:

- In the main window, for example, by selecting **Configuration** > **Firewall** > **Advanced** > **ACL Manager**. In this case, the ACL Manager shows extended ACLs only. These ACLs include those generated by rules you create in the Access Rules, Service Policy Rules, and AAA Rules pages. Be careful that edits you make in ACL Manager do not negatively impact these rules; changes you make here will be reflected on those other pages.
- From a policy that requires an ACL, by clicking the **Manage** button next to the field. In this case, the ACL Manager can have separate tabs for standard and extended ACLs, if the policy allows either type. Otherwise, the view is filtered to show standard, extended, or webtype ACLs only. The ACL Manager never shows EtherType ACLs.

There are separate pages for standard ACLs and webtype ACLs, so that you can configure them in the main window. These pages are functionally equivalent to the ACL Manager without the name:

- Standard ACLs—**Configuration** > **Firewall** > **Advanced** > **Standard ACL**.
- Webtype ACLs—**Configuration** > **Remote Access VPN** > **Clientless SSL VPN Access** > **Advanced** > **Web ACLs**.

ACL Names

Each ACL has a name or numeric ID, such as outside_in, OUTSIDE_IN, or 101. Limit the names to 241 characters or fewer.Consider using all uppercase letters to make it easier to find the name when viewing a running configuration.

Develop a naming convention that will help you identify the intended purpose of the ACL. For example, ASDM uses the convention *interface-name_purpose_direction*, such as “outside_access_in”, for an ACL applied to the “outside” interface in the inbound direction.

Traditionally, ACL IDs were numbers. Standard ACLs were in the range 1-99 or 1300-1999. Extended ACLs were in the range 100-199 or 2000-2699. The ASA does not enforce these ranges, but if you want to use numbers, you might want to stick to these conventions to maintain consistency with routers running IOS Software.

Access Control Entry Order

An ACL is made up of one or more ACEs. Unless you explicitly insert an ACE at a given line, each ACE that you enter for a given ACL name is appended to the end of the ACL.

The order of ACEs is important. When the ASA decides whether to forward or drop a packet, the ASA tests the packet against each ACE in the order in which the entries are listed. After a match is found, no more ACEs are checked.

Thus, if you place a more specific rule after a more general rule, the more specific rule might never be hit. For example, if you want to permit network 10.1.1.0/24, but drop traffic from host 10.1.1.15 on that subnet, the ACE that denies 10.1.1.15 must come before the one that permits 10.1.1.0/24. If the permit 10.1.1.0/24 ACE comes first, 10.1.1.15 will be allowed, and the deny ACE will never be matched.

Use the Up and Down buttons to reposition rules as necessary.

Permit/Deny vs. Match/Do Not Match

Access control entries either “permit” or “deny” traffic that matches the rule. When you apply an ACL to a feature that determines whether traffic is allowed through the ASA or is dropped, such as global and interface access rules, “permit” and “deny” mean what they say.

For other features, such as service policy rules, “permit” and “deny” actually mean “match” or “do not match.” In these cases, the ACL is selecting the traffic that should receive the services of that feature, such as application inspection or redirection to a service module. “Denied” traffic is simply traffic that does not match the ACL, and thus will not receive the service. (In ASDM, service policy rules actually use Match/Do Not Match, and AAA rules use Authenticate/Do Not Authenticate, for example, but in the CLI, it is always permit/deny.)

Access Control Implicit Deny

ACLs that are used for through-the-box access rules have an implicit deny statement at the end. Thus, for traffic controlling ACLs such as those applied to interfaces, if you do not explicitly permit a type of traffic, that traffic is dropped. For example, if you want to allow all users to access a network through the ASA except for one or more particular addresses, then you need to deny those particular addresses and then permit all others.

For management (control plane) ACLs, which control to-the-box traffic, there is no implicit deny at the end of a set of management rules for an interface. Instead, any connection that does not match a management access rule is then evaluated by regular access control rules.

For ACLs used to select traffic for a service, you must explicitly “permit” the traffic; any traffic not “permitted” will not be matched for the service; “denied” traffic bypasses the service.

For EtherType ACLs, the implicit deny at the end of the ACL does not affect IP traffic or ARPs; for example, if you allow EtherType 8037, the implicit deny at the end of the ACL does not now block any IP traffic that you previously allowed with an extended ACL (or implicitly allowed from a high security interface to a low security interface). However, if you *explicitly* deny all traffic with an EtherType ACE, then IP and ARP traffic is denied; only physical protocol traffic, such as auto-negotiation, is still allowed.

IP Addresses Used for Extended ACLs When You Use NAT

When you use NAT or PAT, you are translating addresses or ports, typically mapping between internal and external addresses. If you need to create an extended ACL that applies to addresses or ports that have been translated, you need to determine whether to use the real (untranslated) addresses or ports or the mapped ones. The requirement differs by feature.

Using the real address and port means that if the NAT configuration changes, you do not need to change the ACLs.

Features That Use Real IP Addresses

The following commands and features use real IP addresses in the ACLs, even if the address as seen on an interface is the mapped address:

- Access Rules (extended ACLs referenced by the **access-group** command)
- Service Policy Rules (Modular Policy Framework **match access-list** command)
- Botnet Traffic Filter traffic classification (**dynamic-filter enable classify-list** command)

- AAA Rules (**aaa ... match** commands)
- WCCP (**wccp redirect-list group-list** command)

For example, if you configure NAT for an inside server, 10.1.1.5, so that it has a publicly routable IP address on the outside, 209.165.201.5, then the access rule to allow the outside traffic to access the inside server needs to reference the server's real IP address (10.1.1.5), and not the mapped address (209.165.201.5).

Features That Use Mapped IP Addresses

The following features use ACLs, but these ACLs use the mapped values as seen on an interface:

- IPsec ACLs
- **capture** command ACLs
- Per-user ACLs
- Routing protocol ACLs
- All other feature ACLs.

Time-Based ACEs

You can apply time range objects to extended and webtype ACEs so that the rules are active for specific time periods only. These types of rules let you differentiate between activity that is acceptable at certain times of the day but that is unacceptable at other times. For example, you could provide additional restrictions during working hours, and relax them after work hours or at lunch. Conversely, you could essentially shut your network down during non-work hours.

You cannot create time-based rules that have the exact same protocol, source, destination, and service criteria of a rule that does not include a time range object. The non-time-based rule always overrides the duplicate time-based rule, as they are redundant.



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 ASA finishes any currently running task and then services the command to deactivate the ACL.

Licensing for Access Control Lists

Access control lists do not require a special license.

However, to use **sctp** as the protocol in an entry, you must have a Carrier license.

Guidelines for ACLs

Firewall Mode

- Extended and standard ACLs are supported in routed and transparent firewall modes.
- Webtype ACLs are supported in routed mode only.
- EtherType ACLs are supported for bridge group member interfaces only, in routed and transparent modes.

Failover and Clustering

Configuration sessions are not synchronized across failover or clustered units. When you commit the changes in a session, they are made in all failover and cluster units as normal.

IPv6

- Extended and webtype ACLs allow a mix of IPv4 and IPv6 addresses.
- Standard ACLs do not allow IPv6 addresses.
- EtherType ACLs do not contain IP addresses.

Additional Guidelines

- When you specify a network mask, the method is different from the Cisco IOS software **access-list** command. The ASA uses a network mask (for example, 255.255.255.0 for a Class C mask). The Cisco IOS mask uses wildcard bits (for example, 0.0.0.255).
- Normally, you cannot reference an object or object group that does not exist in an ACL or object group, or delete one that is currently referenced. You also cannot reference an ACL that does not exist in an **access-group** command (to apply access rules). However, you can change this default behavior so that you can "forward reference" objects or ACLs before you create them. Until you create the objects or ACLs, any rules or access groups that reference them are ignored. To enable forward referencing, select the option in the access rules advanced settings; choose **Configuration > Access Rules** and click the **Advanced** button.
- If you enter more than one item in source or destination address, or source or destination service, ASDM automatically creates an object group for them with the prefix DM_INLINE. These objects are automatically expanded to their component parts in the rule table view, but you can see the object names if you deselect the **Auto-expand network and service objects with specified prefix** rule table preference in **Tools** > **Preferences**.
- (Extended ACL only) The following features use ACLs, but cannot accept an ACL with identity firewall (specifying user or group names), FQDN (fully-qualified domain names), or Cisco TrustSec values:
  - VPN **crypto map** command
  - VPN **group-policy** command, except for **vpn-filter**
  - WCCP
  - DAP

Configure ACLs

The following sections explain how to configure the various types of generic ACL, except those used as access rules (including EtherType), service policy rules, AAA rules, and other uses where ASDM provides a special-purpose page for those rule-based policies.

Configure Extended ACLs

An extended ACL is represented as a named container of ACEs. To create a new ACL, you must first create the container. Then, you can add ACEs, edit existing ACEs, and reorder the ACEs using the table in ACL Manager.

The extended ACL can include a mix of IPv4 and IPv6 addresses.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Advanced** > **ACL Manager**.

Step 2 If you are creating a new ACL, choose **Add** > **Add ACL**, fill in a name, and click **OK**.

The ACL container is added to the table. You can later rename it by selecting it and clicking **Edit**.

Step 3 Do any of the following:

- To add an ACE at the end of the ACL, select the ACL name or any ACE within it and choose **Add > Add ACE**.
- To insert an ACE at a specific location, select an existing ACE and choose **Add > Insert** to add the ACE above the rule, or choose **Add > Insert After**.
- To edit a rule, select it and click **Edit**.

Step 4 Fill in the ACE properties. The primary options to select are:

- **Action: Permit/Deny**—Whether you are permitting (selecting) the described traffic or are denying (deselecting, not matching) it.
- **Source/Destination criteria**—A definition of the source (originating address) and destination (target address of the traffic flow). You typically configure IPv4 or IPv6 addresses of hosts or subnets, which you can represent with network or network object groups. You can also specify a user or user group name for the source. Additionally, you can use the Service field to identify the specific type of traffic if you want to focus the rule more narrowly than all IP traffic. If you implement Cisco TrustSec, you can use security groups to define source and destination.

For detailed information on all of the available options, see Extended ACE Properties, on page 48.

When you are finished defining the ACE, click **OK** to add the rule to the table.

Step 5 Click **Apply**.

Extended ACE Properties

When you add or edit an ACE in an extended ACL, you can configure the following properties. In many fields, you can click the "..." button on the right of the edit box to select, create, or edit objects that are available for the field.

Action: Permit/Deny

Whether you are permitting (selecting) the described traffic or are denying (deselecting, not matching) it.

Source Criteria

The characteristics of the originator of the traffic you are trying to match. You must configure Source, but the other properties are optional.

Source

The IPv4 or IPv6 address of the source. The default is **any**, which matches all IPv4 or IPv6 addresses; you can use **any4** to target IPv4 only, or **any6** to target IPv6 only. You can specify a single host address (such as 10.100.10.5 or 2001:DB8::0DB8:800:200C:417A), a subnet (in 10.100.10.0/24 or 10.100.10.0/255.255.255.0 format, or for IPv6, 2001:DB8:0:CD30::/60), the name of a network object or network object group, or the name of an interface.

User

If you enable the identity firewall, you can specify a user or user group as the traffic source. The IP address the user is currently using will match the rule. You can specify a username (DOMAIN\user), a user group (DOMAIN\\group, note the double \ indicates a group name), or a user object group. For this field, it is far easier to click "**...**" to select names from your AAA server group than to type them in.

Security Group

If you enable Cisco TrustSec, you can specify a security group name or tag (1-65533), or security group object.

More Options > Source Service

If you specify TCP, UDP, or SCTP as the destination service, you can optionally specify a predefined service object for TCP, UDP, TCP-UDP, or SCTP, or use your own object. Typically, you define the destination service only and not the source service. Note that if you define the source service, the destination service protocol must match it (for example, both TCP, with or without port definitions).

Destination Criteria

The characteristics of the target of the traffic you are trying to match. You must configure Destination, but the other properties are optional.

Destination

The IPv4 or IPv6 address of the destination. The default is **any**, which matches all IPv4 or IPv6 addresses; you can use **any4** to target IPv4 only, or **any6** to target IPv6 only. You can specify a single host address (such as 10.100.10.5 or 2001:DB8::0DB8:800:200C:417A), a subnet (in 10.100.10.0/24 or 10.100.10.0/255.255.255.0 format, or for IPv6, 2001:DB8:0:CD30::/60), the name of a network object or network object group, or the name of an interface.

Security Group

If you enable Cisco TrustSec, you can specify a security group name or tag (1-65533), or security group object.

Service

The protocol of the traffic, such as IP, TCP, UDP, and optionally ports for TCP, UDP, or SCTP. The default is IP, but you can select a more specific protocol to target traffic with more granularity. Typically, you would select some type of service object. For TCP, UDP, and SCTP, you can specify ports, for example, tcp/80, tcp/http, tcp/10-20 (for a range of ports), tcp-udp/80 (match any TCP or UDP traffic on port 80), sctp/diameter, and so forth. For detailed information on specifying services, see Service Specifications in Extended ACEs, on page 50.

Description

A explanation of the purpose of the ACE, up to 100 characters per line. You can enter multiple lines; each line is added as a remark in the CLI, and the remarks are placed before the ACE.

Note If you add remarks with non-English characters on one platform (such as Windows) then try to remove them from another platform (such as Linux), you might not be able to edit or delete them because the original characters might not be correctly recognized. This limitation is due to an underlying platform dependency that encodes different language characters in different ways.

Enable Logging; Logging Level; More Options > Logging Interval

The logging options define how syslog messages will be generated for rules. These options apply to ACLs that are used as access rules only, that is, those attached to interfaces or applied globally. The options are ignored for ACLs used for other features. You can implement the following logging options:

Deselect Enable Logging

This will disable logging for the rule. No syslog messages of any type will be issued for connections that match this rule.

Select Enable Logging with Logging Level = Default

This provides the default logging for rules. Syslog message 106023 is issued for each denied connection. If the appliance comes under attack, the frequency of issuing this message could impact services.

Select Enable Logging with Non-Default Logging Level

This provides a summarized syslog message, 106100, instead of 106023. Message 106100 is issued upon first hit, then again at each interval configured in **More Options > Logging Interval** (default is every 300 seconds, you can specify 1-600), showing the number of hits during that interval. The recommended logging level is **Informational**.

Summarizing deny messages can reduce the impact of attacks and possibly make it easier for you to analyze messages. If you do come under a denial of service attack, you might see message 106101, which indicates that the number of cached deny flows used to produce the hit count for message 106100 has exceeded the maximum for an interval. At this point, the appliance stops collecting statistics until the next interval to mitigate the attack.

More Options > Enable Rule

Whether the rule is active on the device. Disabled rules appear with strike-through text in the rule table. Disabling a rule lets you stop its application to traffic without deleting it, so you can enable it again later if you decide you need it.

More Options > Time Range

The name of the time range object that defines the times of day and days of the week when the rule should be active. If you do not specify a time range, the rule is always active.

Service Specifications in Extended ACEs

For the destination service in an extended ACE, you can specify any of the following criteria. The options are similar, but more limited, for source service, which is limited to TCP, UDP, TCP-UDP, or SCTP criteria.

Object name

The name of any type of service object or service object group. These objects can include many of the specifications explained below, allowing you to easily reuse service definitions among ACLs. There are many pre-defined objects, so you might find what you need without having to manually type the specification or create your own objects.

Protocol

A number between 1-255, or a well-known name, such as **ip**, **tcp**, **udp**, **gre**, and so forth.

TCP, UDP, TCP-UDP, SCTP ports

You can include port specifications on the **tcp**, **udp**, **tcp-udp**, and **sctp** keywords. The tcp-udp keyword lets you define ports for both protocols without having to specify them separately. You can use the following methods to specify ports:

- Single port—tcp/80, udp/80, tcp-udp/80, sctp/3868, or a well-known service name, such as tcp/www, udp/snmp, or sctp/diameter.
- Range of ports—tcp/1-100, udp/1-100, tcp-udp/1-100, sctp/1-100 matches ports 1-100 inclusive.
- Not equal to a port—Add != to the beginning of the specification, for example, !=tcp/80 to match any TCP traffic except TCP port 80 (HTTP).
- Less than a port number—Add <, for example <tcp/150 to match TCP traffic for any port below 150.
- Greater than a port number—Add >, for example, >tcp150 to match TCP traffic for any port above 150.



Note 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.

ICMP, ICMP6 messages

You can target specific messages (such as ping echo request and reply messages) and even message codes. There are many pre-defined objects that cover ICMP (for IPv4) and ICMP6 (for IPv6), so you might not need to manually define the criteria. The format is:

icmp/*icmp_message_type*[/*icmp_message_code*]

icmp6/*icmp6_message_type*[/*icmp6_message_code*]

Where the message type is 1-255 or a well-known name, and the code is 0-255. Ensure that the number you select matches to an actual type/code or the ACE will never be matched.

Configure Standard ACLs

A standard ACL is represented as a named container of ACEs. To create a new ACL, you must first create the container. Then, you can add ACEs, edit existing ACEs, and reorder the ACEs using the standard ACL table. The table can appear as a tab in the ACL Manager when you configure ACLs while configuring the policies that use them, in which case the procedures are the same except for how you get to the window.

A standard ACL uses IPv4 addresses only, and defines destination addresses only.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Advanced** > **Standard ACL**.

Step 2 If you are creating a new ACL, choose **Add** > **Add ACL**, fill in a name, and click **OK**.

The ACL container is added to the table. You cannot rename a standard ACL.

Step 3 Do any of the following:

- To add an ACE at the end of the ACL, select the ACL name or any ACE within it and choose **Add > Add ACE**.
- To insert an ACE at a specific location, select an existing ACE and choose **Add > Insert** to add the ACE above the rule, or choose **Add > Insert After**.
- To edit a rule, select it and click **Edit**.

Step 4 Fill in the ACE properties. The options are:

- **Action: Permit/Deny**—Whether you are permitting (selecting) the described traffic or are denying (deselecting, not matching) it.
- **Address**—A definition of the destination or target address of the traffic flow. You can specify a host address such as 10.100.1.1, a network (in 10.100.1.0/24 or 10.100.1.0/255.255.255.0 format), or you can select a network object (which simply loads the contents of the object into the Address field).
- **Description**—A explanation of the purpose of the ACE, up to 100 characters per line. You can enter multiple lines; each line is added as a remark in the CLI, and the remarks are placed before the ACE.

  Note
  If you add remarks with non-English characters on one platform (such as Windows) then try to remove them from another platform (such as Linux), you might not be able to edit or delete them because the original characters might not be correctly recognized. This limitation is due to an underlying platform dependency that encodes different language characters in different ways.

When you are finished defining the ACE, click **OK** to add the rule to the table.

Step 5 Click **Apply**.

Configure Webtype ACLs

Webtype ACLs are used for filtering clientless SSL VPN traffic, constraining user access to specific networks, subnets, hosts, and Web servers. If you do not define a filter, all connections are allowed. A webtype ACL is represented as a named container of ACEs. To create a new ACL, you must first create the container. Then, you can add ACEs, edit existing ACEs, and reorder the ACEs using the Web ACL table. The table appears as the ACL Manager when you configure webtype ACLs while configuring the policies that use them, in which case the procedures are the same except for how you get to the window.

The webtype ACL can include a mix of IPv4 and IPv6 addresses in addition to URL specifications.

Procedure

Step 1 Choose **Configuration** > **Remote Access VPN** > **Clientless SSL VPN Access** > **Advanced** > **Web** > **ACLs**.

Step 2 If you are creating a new ACL, choose **Add** > **Add ACL**, fill in a name, and click **OK**.

The ACL container is added to the table. You can later rename it by selecting it and clicking **Edit**.

Step 3 Do any of the following:

- To add an ACE at the end of the ACL, select the ACL name or any ACE within it and choose **Add > Add ACE**.
- To insert an ACE at a specific location, select an existing ACE and choose **Add > Insert** to add the ACE above the rule, or choose **Add > Insert After**.
- To edit a rule, select it and click **Edit**.

Step 4 Fill in the ACE properties. The primary options to select are:

- **Action: Permit/Deny**—Whether you are permitting (selecting) the described traffic or are denying (deselecting, not matching) it.
- **Filter**—The traffic matching criteria, based on the destination. You can either specify a URL by selecting the protocol and entering the server name and optionally, path and file name, or you can specify a destination IPv4 or IPv6 address and TCP service.

For detailed information on all of the available options, see Webtype ACE Properties, on page 53.

When you are finished defining the ACE, click **OK** to add the rule to the table.

Step 5 Click **Apply**.

Webtype ACE Properties

When you add or edit an ACE in a webtype ACL, you can configure the following properties. In many fields, you can click the "..." button on the right of the edit box to select, create, or edit objects that are available for the field.

For a given ACE, you can filter on URL or Address, but not both.

- **Action: Permit/Deny**—Whether you are permitting (selecting) the described traffic or are denying (deselecting, not matching) it.
  - **Filter on URL**—Match traffic based on destination URL. Select the protocol and enter the server name and optionally, path and file name. For example, http://www.example.com or to cover all servers, http://*.example.com. Following are some tips and limitations on specifying URLs:
    - Select **any** to match all URLs.
    - 'Permit url any' will allow all the URLs that have the format protocol://server-ip/path and will block traffic that does not match this pattern, such as port-forwarding. There should be an ACE to allow connections to the required port (port 1494 in the case of Citrix) so that an implicit deny does not occur.
    - Smart tunnel and ica plug-ins are not affected by an ACL with 'permit url any' because they match smart-tunnel:// and ica:// types only.
    - You can use these protocols: cifs://, citrix://, citrixs://, ftp://, http://, https://, imap4://, nfs://, pop3://, smart-tunnel://, and smtp://. You can also use wildcards in the protocol; for example, htt* matches http and https, and an asterisk * matches all protocols. For example, *://*.example.com matches any type URL-based traffic to the example.com network.
    - If you specify a smart-tunnel:// URL, you can include the server name only. The URL cannot contain a path. For example, smart-tunnel://www.example.com is acceptable, but smart-tunnel://www.example.com/index.html is not.
    - An asterisk * matches none or any number of characters. To match any http URL, enter http://*/*.
    - A question mark ? matches any one character exactly.
    - Square brackets [] are range operators, matching any character in the range. For example, to match both http://www.cisco.com:80/ and http://www.cisco.com:81/, enter **http://www.cisco.com:8[01]/**.
  - **Filter on Address and Service**—Match traffic based on destination address and service.
    - **Address**—The IPv4 or IPv6 address of the destination. To match all addresses, you can use **any**, which matches all IPv4 or IPv6 addresses, **any4** to match IPv4 only, or **any6** to match IPv6 only. You can specify a single host address (such as 10.100.10.5 or 2001:DB8::0DB8:800:200C:417A), a subnet (in 10.100.10.0/24 or 10.100.10.0/255.255.255.0 format, or for IPv6, 2001:DB8:0:CD30::/60), or select a network object, which fills in the field with the contents of the object.
    - **Service**—A single TCP service specification. The default is **tcp** with no ports, but you can specify a single port (such as tcp/80 or tcp/www) or port range (such as tcp/1-100). You can include operators; for example, !=tcp/80 excludes port 80; <tcp/80 is all ports less than 80; >tcp/80 is all ports greater than 80.

- **Enable Logging**; **Logging Level**; **More Options > Logging Interval**—The logging options define how syslog messages will be generated for rules that actually deny traffic. You can implement the following logging options:
  - **Deselect Enable Logging**—This will disable logging for the rule. No syslog messages of any type will be issued for traffic denied by this rule.
  - **Select Enable Logging with Logging Level = Default**—This provides the default logging for rules. Syslog message 106103 is issued for each denied packet. If the appliance comes under attack, the frequency of issuing this message could impact services.
  - **Select Enable Logging with Non-Default Logging Level**—This provides a summarized syslog message, 106102, instead of 106103. Message 106102 is issued upon first hit, then again at each interval configured in **More Options > Logging Interval** (default is every 300 seconds, you can specify 1-600), showing the number of hits during that interval. The recommended logging level is **Informational**.
- **More Options > Time Range**—The name of the time range object that defines the times of day and days of the week when the rule should be active. If you do not specify a time range, the rule is always active.

Examples for Webtype ACLs

Following are some examples of URL-based rules for webtype ACLs.

	Filter	Effect
Deny	url http://*.yahoo.com/	Denies access to all of Yahoo!
Deny	url cifs://fileserver/share/directory	Denies access to all files in the specified location.
Deny	url https://www.example.com/ directory/file.html	Denies access to the specified file.
Permit	url https://www.example.com/directory	Permits access to the specified location.
Deny	url http://*:8080/	Denies HTTPS access to anywhere via port 8080.
Deny	url http://10.10.10.10	Denies HTTP access to 10.10.10.10.
Permit	url any	Permits access to any URL. Usually used after an ACL that denies url access.

Monitoring ACLs

The ACL Manager, Standard ACL, Web ACL, and EtherType ACL tables show a consolidated view of ACLs. But to see exactly what is configured on the device, you can use the following commands. Choose **Tools > Command Line Interface** to enter the commands.

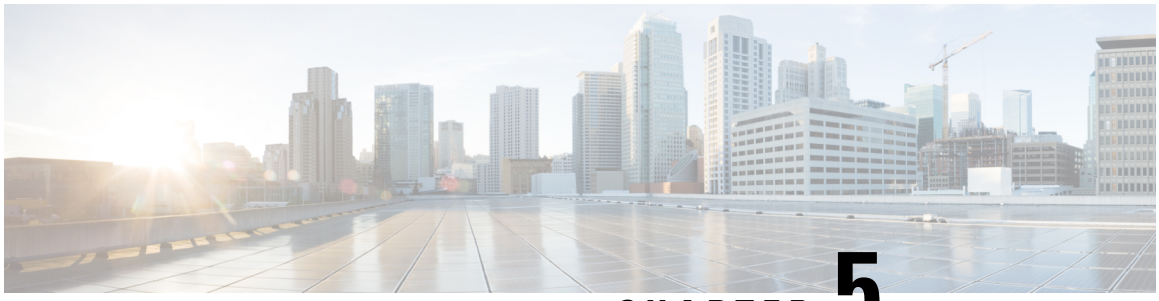
- **show access-list** [*name*]—Displays the access lists, including the line number for each ACE and hit counts. Include an ACL name or you will see all access lists.

- **show running-config access-list** [*name*]—Displays the current running access-list configuration. Include an ACL name or you will see all access lists.

History for ACLs

Feature Name	Releases	Description
Extended, standard, webtype ACLs	7.0(1)	ACLs are used to control network access or to specify traffic for many features to act upon. An extended access control list is used for through-the-box access control and several other features. Standard ACLs are used in route maps and VPN filters. Webtype ACLs are used in clientless SSL VPN filtering. EtherType ACLs control non-IP layer 2 traffic. We added the ACL Manager and other pages for configuring ACLs.
Real IP addresses in extended ACLs	8.3(1)	When using NAT or PAT, mapped addresses and ports are no longer used in an ACL for several features. You must use the real, untranslated addresses and ports for these features. Using the real address and port means that if the NAT configuration changes, you do not need to change the ACLs.
Support for Identity Firewall in extended ACLs	8.4(2)	You can now use identity firewall users and groups for the source and destination. You can use an identity firewall ACL with access rules, AAA rules, and for VPN authentication.
EtherType ACL support for IS-IS traffic	8.4(5), 9.1(2)	In transparent firewall mode, the ASA can now control IS-IS traffic using an EtherType ACL. We modified the following screen: Configuration > Device Management > Management Access > EtherType Rules.
Support for Cisco TrustSec in extended ACLs	9.0(1)	You can now use Cisco TrustSec security groups for the source and destination. You can use an identity firewall ACL with access rules.
Unified extended and webtype ACLs for IPv4 and IPv6	9.0(1)	Extended and webtype ACLs now support IPv4 and IPv6 addresses. You can even specify a mix of IPv4 and IPv6 addresses for the source and destination. The **any** keyword was changed to represent IPv4 and IPv6 traffic. The **any4** and **any6** keywords were added to represent IPv4-only and IPv6-only traffic, respectively. The IPv6-specific ACLs are deprecated. Existing IPv6 ACLs are migrated to extended ACLs. See the release notes for more information about migration. We modified the following screens: Configuration > Firewall > Access Rules Configuration > Remote Access VPN > Network (Client) Access > Group Policies > General > More Options

Feature Name	Releases	Description
Extended ACL and object enhancement to filter ICMP traffic by ICMP code	9.0(1)	ICMP traffic can now be permitted/denied based on ICMP code. We introduced or modified the following screens: Configuration > Firewall > Objects > Service Objects/Groups Configuration > Firewall > Access Rule
Configuration session for editing ACLs and objects. Forward referencing of objects and ACLs in access rules.	9.3(2)	You can now edit ACLs and objects in an isolated configuration session. You can also forward reference objects and ACLs, that is, configure rules and access groups for objects or ACLs that do not yet exist. We modified the Advanced settings for access rules.
ACL support for Stream Control Transmission Protocol (SCTP)	9.5(2)	You can now create ACL rules using the **sctp** protocol, including port specifications. We modified the add/edit dialog boxes for access control entries on the **Configuration** > **Firewall** > **Advanced** > **ACL Manager** page.
Ethertype rule support for the IEEE 802.2 Logical Link Control packet's Destination Service Access Point address.	9.6(2)	You can now write Ethertype access control rules for the IEEE 802.2 Logical Link Control packet's Destination Service Access Point address. Because of this addition, the **bpdu** keyword no longer matches the intended traffic. Rewrite **bpdu** rules for **dsap 0x42**. We modified the following screen: **Configuration** > **Firewall** > **EtherType Rules**.
Support in routed mode for Ethertype rules on bridge group member interfaces and extended access rules on Bridge Group Virtual Interfaces (BVI).	9.7(1)	You can now create Ethertype ACLs and apply them to bridge group member interfaces in routed mode. You can also apply extended access rules to the Bridge Virtual Interface (BVI) in addition to the member interfaces. We modified the following screens: **Configuration** > **Firewall** > **Access Rules**, **Configuration** > **Firewall** > **Ethertype Rules**.
EtherType access control list changes.	9.9(1)	EtherType access control lists now support Ethernet II IPX (EII IPX). In addition, new keywords are added to the DSAP keyword to support common DSAP values: BPDU (0x42), IPX (0xE0), Raw IPX (0xFF), and ISIS (0xFE). Consequently, existing EtherType access control entries that use the BPDU or ISIS keywords will be converted automatically to use the DSAP specification, and rules for IPX will be converted to 3 rules (DSAP IPX, DSAP Raw IPX, and EII IPX). In addition, packet capture that uses IPX as an EtherType value has been deprecated, because IPX corresponds to 3 separate EtherTypes. We modified the following screens: **Configuration** > **Firewall** > **Ethertype Rules**.



CHAPTER 5

Identity Firewall

This chapter describes how to configure the ASA for the Identity Firewall.

- About the Identity Firewall, on page 57
- Guidelines for the Identity Firewall, on page 63
- Prerequisites for the Identity Firewall, on page 65
- Configure the Identity Firewall, on page 66
- Monitoring the Identity Firewall, on page 72
- History for the Identity Firewall, on page 72

About the Identity Firewall

In an enterprise, users often need access to one or more server resources. Typically, a firewall is not aware of the users' identities and, therefore, cannot apply security policies based on identity. To configure per-user access policies, you must configure a user authentication proxy, which requires user interaction (a username/password query).

The Identity Firewall in the ASA provides more granular access control based on users' identities. You can configure access rules and security policies based on user names and user group names rather than through source IP addresses. The ASA applies the security policies based on an association of IP addresses to Windows Active Directory login information and reports events based on the mapped usernames instead of network IP addresses.

The Identity Firewall integrates with Microsoft Active Directory in conjunction with an external Active Directory (AD) Agent that provides the actual identity mapping. The ASA uses Windows Active Directory as the source to retrieve the current user identity information for specific IP addresses and allows transparent authentication for Active Directory users.

Identity-based firewall services enhance the existing access control and security policy mechanisms by allowing users or groups to be specified in place of source IP addresses. Identity-based security policies can be interleaved without restriction between traditional IP address-based rules.

The key benefits of the Identity Firewall include:

- Decoupling network topology from security policies
- Simplifying the creation of security policies
- Providing the ability to easily identify user activities on network resources
- Simplifying user activity monitoring

Architecture for Identity Firewall Deployments

The Identity Firewall integrates with Window Active Directory in conjunction with an external Active Directory (AD) Agent that provides the actual identity mapping.

The identity firewall consists of three components:

- ASA
- Microsoft Active Directory

  Although Active Directory is part of the Identity Firewall on the ASA, Active Directory administrators manage it. The reliability and accuracy of the data depends on data in Active Directory.

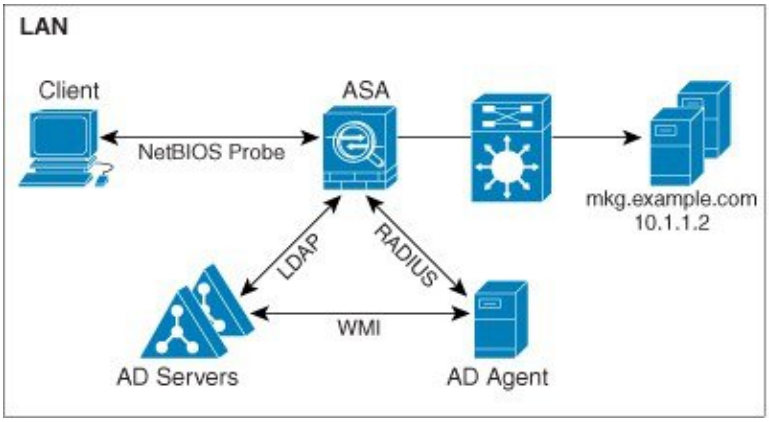
  Supported versions include Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 servers.
- Active Directory (AD) Agent

  The AD Agent runs on a Windows server. Supported Windows servers include Windows 2003, Windows 2008, and Windows 2008 R2.

Note Windows 2003 R2 is not supported for the AD Agent server.

The following figure show the components of the Identity Firewall. The succeeding table describes the roles of these components and how they communicate with one another.

Figure 3: Identity Firewall Components



1	**On the ASA**: Administrators configure local user groups and Identity Firewall policies.	4	**Client <-> ASA**: The client logs into the network through Microsoft Active Directory. The AD Server authenticates users and generates user login security logs. Alternatively, the client can log into the network through a cut-through proxy or VPN.

2	**ASA <-> AD Server**: The ASA sends an LDAP query for the Active Directory groups configured on the AD Server. The ASA consolidates local and Active Directory groups and applies access rules and Modular Policy Framework security policies based on user identity.	5	**ASA <-> Client**: Based on the policies configured on the ASA, it grants or denies access to the client. If configured, the ASA probes the NetBIOS of the client to pass inactive and no-response users.
3	**ASA <-> AD Agent**: Depending on the Identity Firewall configuration, the ASA downloads the IP-user database or sends a RADIUS request to the AD Agent that asks for the user's IP address. The ASA forwards the new mapped entries that have been learned from web authentication and VPN sessions to the AD Agent.	6	**AD Agent <-> AD Server**: The AD Agent maintains a cache of user ID and IP address mapped entries. and notifies the ASA of changes. The AD Agent sends logs to a syslog server.

Features of the Identity Firewall

The Identity Firewall includes the following key features.

Flexibility

- The ASA can retrieve user identity and IP address mapping from the AD Agent by querying the AD Agent for each new IP address or by maintaining a local copy of the entire user identity and IP address database.
- Supports host group, subnet, or IP address for the destination of a user identity policy.
- Supports a fully qualified domain name (FQDN) for the source and destination of a user identity policy.
- Supports the combination of 5-tuple policies with ID-based policies. The identity-based feature works in tandem with the existing 5-tuple solution.
- Supports use with application inspection policies.
- Retrieves user identity information from remote access VPN, AnyConnect VPN, L2TP VPN and cut-through proxy. All retrieved users are populated to all ASAs that are connected to the AD Agent.

Scalability

- Each AD Agent supports 100 ASAs. Multiple ASAs are able to communicate with a single AD Agent to provide scalability in larger network deployments.
- Supports 30 Active Directory servers provided the IP address is unique among all domains.
- Each user identity in a domain can have up to 8 IP addresses.
- Supports up to 64,000 user identity-IP address mapped entries in active policies for the ASA 5500 Series models. This limit controls the maximum number of users who have policies applied. The total number of users are the aggregate of all users configured in all different contexts.
- Supports up to 512 user groups in active ASA policies.

- A single access rule can contain one or more user groups or users.
- Supports multiple domains.

Availability

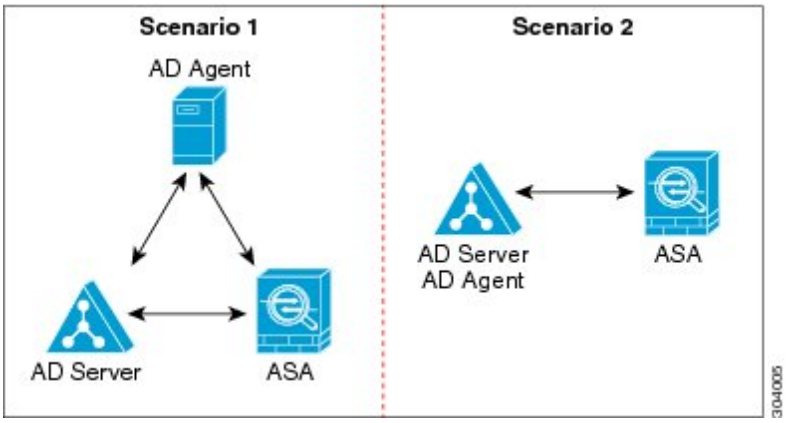
- The ASA retrieves group information from the Active Directory and falls back to web authentication for IP addresses when the AD Agent cannot map a source IP address to a user identity.
- The AD Agent continues to function when any of the Active Directory servers or the ASA are not responding.
- Supports configuring a primary AD Agent and a secondary AD Agent on the ASA. If the primary AD Agent stops responding, the ASA can switch to the secondary AD Agent.
- If the AD Agent is unavailable, the ASA can fall back to existing identity sources such as cut-through proxy and VPN authentication.
- The AD Agent runs a watchdog process that automatically restarts its services when they are down.
- Allows a distributed IP address/user mapping database for use among ASAs.

Deployment Scenarios

You can deploy the components of the Identity Firewall in the following ways, depending on your environmental requirements.

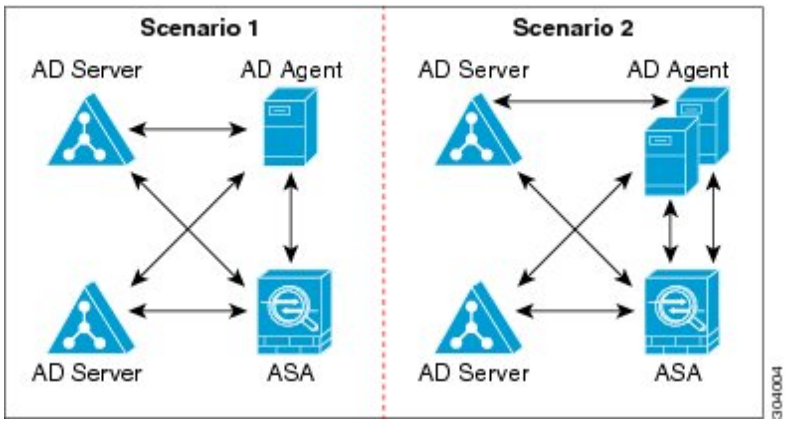
The following figure shows how you can deploy the components of the Identity Firewall to allow for redundancy. Scenario 1 shows a simple installation without component redundancy. Scenario 2 also shows a simple installation without redundancy. However, in this deployment scenario, the Active Directory server and AD Agent are co-located on the same Windows server.

Figure 4: Deployment Scenario without Redundancy



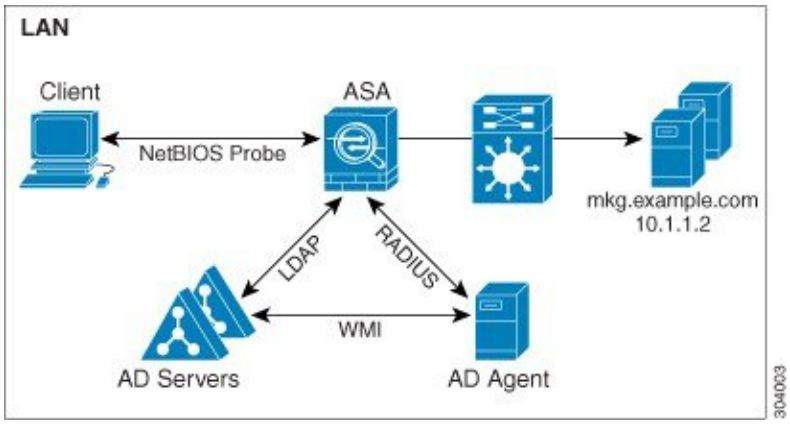
The following figure shows how you can deploy the Identity Firewall components to support redundancy. Scenario 1 shows a deployment with multiple Active Directory servers and a single AD Agent installed on a separate Windows server. Scenario 2 shows a deployment with multiple Active Directory servers and multiple AD Agents installed on separate Windows servers.

Figure 5: Deployment Scenario with Redundant Components



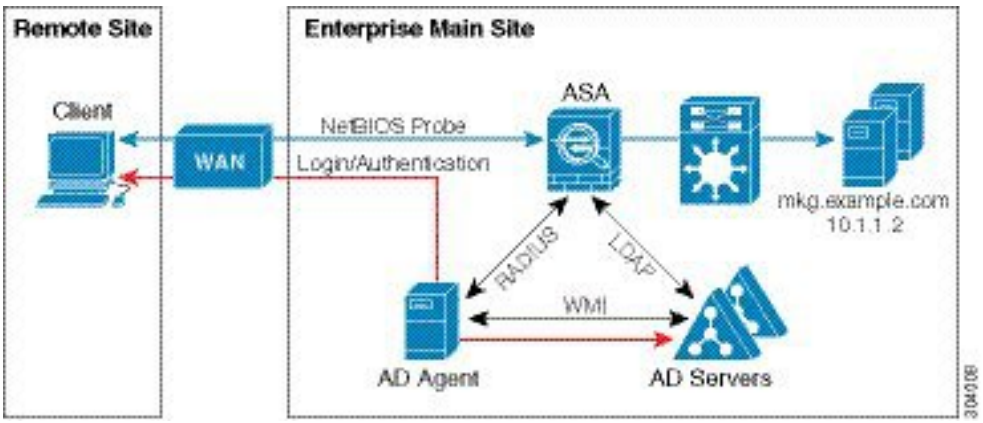
The following figure shows how all Identity Firewall components—Active Directory server, the AD Agent, and the clients—are installed and communicate on the LAN.

Figure 6: LAN -based Deployment



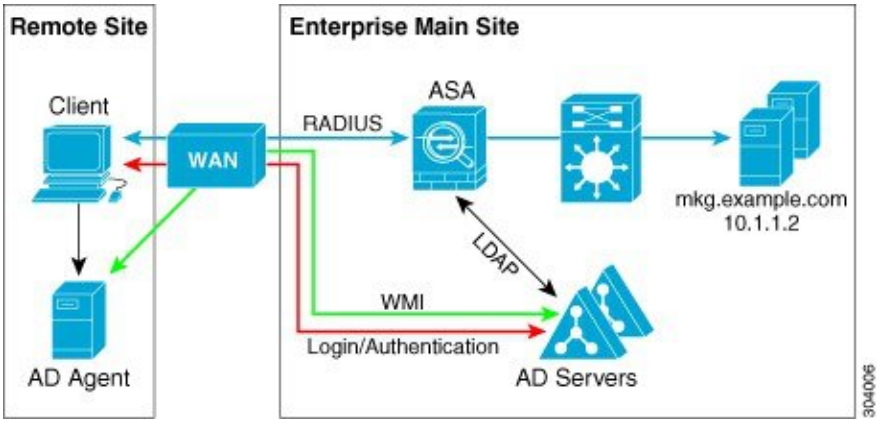
The following figure shows a WAN-based deployment to support a remote site. The Active Directory server and the AD Agent are installed on the main site LAN. The clients are located at a remote site and connect to the Identity Firewall components over a WAN.

Figure 7: WAN-based Deployment



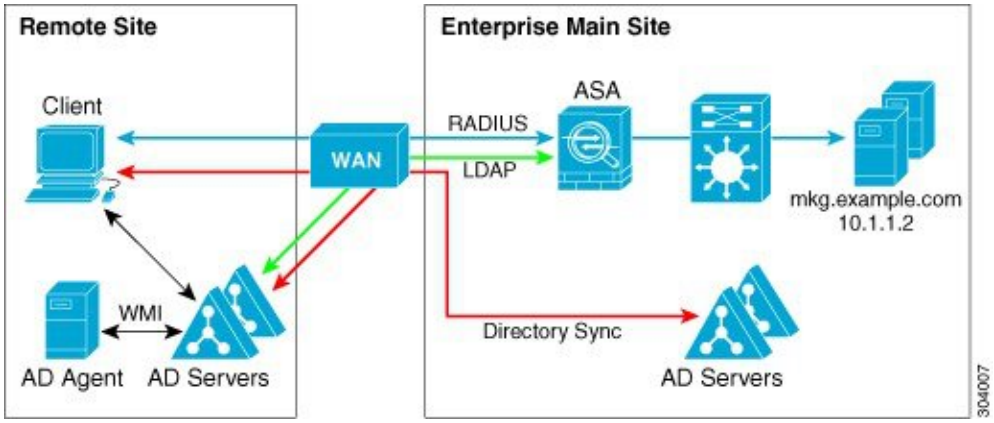
The following figure also shows a WAN-based deployment to support a remote site. The Active Directory server is installed on the main site LAN. However, the AD Agent is installed and accessed by the clients at the remote site. The remote clients connect to the Active Directory servers at the main site over a WAN.

Figure 8: WAN-based Deployment with Remote AD Agent



The following figure shows an expanded remote site installation. An AD Agent and Active Directory servers are installed at the remote site. The clients access these components locally when logging into network resources located at the main site. The remote Active Directory server must synchronize its data with the central Active Directory servers located at the main site.

Figure 9: WAN-based Deployment with Remote AD Agent and AD Servers



Guidelines for the Identity Firewall

This section describes the guidelines and limitations that you should check before configuring the Identity Firewall.

Failover

- The Identity Firewall supports user identity-IP address mapping and AD Agent status replication from active to standby when Stateful Failover is enabled. However, only user identity-IP address mapping, AD Agent status, and domain status are replicated. User and user group records are not replicated to the standby ASA.
- When failover is configured, the standby ASA must also be configured to connect to the AD Agent directly to retrieve user groups. The standby ASA does not send NetBIOS packets to clients even when the NetBIOS probing options are configured for the Identity Firewall.
- When a client is determined to be inactive by the active ASA, the information is propagated to the standby ASA. User statistics are not propagated to the standby ASA.
- When you have failover configured, you must configure the AD Agent to communicate with both the active and standby ASAs. See the *Installation and Setup Guide for the Active Directory Agent* for the steps to configure the ASA on the AD Agent server.

IPv6

- The AD Agent supports endpoints with IPv6 addresses. It can receive IPv6 addresses in log events, maintain them in its cache, and send them through RADIUS messages. The AAA server must use an IPv4 address.
- NetBIOS over IPv6 is not supported.

Additional Guidelines

- A full URL as a destination address is not supported.

- For NetBIOS probing to function, the network between the ASA, AD Agent, and clients must support UDP-encapsulated NetBIOS traffic.
- MAC address checking by the Identity Firewall does not work when intervening routers are present. Users logged into clients that are behind the same router have the same MAC addresses. With this implementation, all the packets from the same router are able to pass the check, because the ASA is unable to ascertain the actual MAC addresses behind the router.
- Although you can use user specifications in VPN filter ACLs, the user-based rules are interpreted uni-directionally rather than bi-directionally, which is how VPN filter usually works. That is, you can filter based on user-initiated traffic, but the filter does not apply for going from the destination back to the user. For example, you could include a rule that allows a specific user to ping a server, but that rule will not allow the server to ping the user.
- The following ASA features do not support using the identity-based object and FQDN in an extended ACL:
  - Crypto maps
  - WCCP
  - NAT
  - Group policy (except for VPN filters)
  - DAP
- You can use the **user-identity update active-user-database** command to actively initiate a user-IP address download from the AD agent.

  By design, if a previous download session has finished, the ASA does not allow you to issue this command again.

  As a result, if the user-IP database is very large, the previous download session is not finished yet, and you issue another **user-identity update active-user-database** command, the following error message appears:

  ```
  "ERROR: one update active-user-database is already in progress."
  ```

  You need to wait until the previous session is completely finished, then you can issue another **user-identity update active-user-database** command.

  Another example of this behavior occurs because of packet loss from the AD Agent to the ASA.

  When you issue a **user-identity update active-user-database** command, the ASA requests the total number of user-IP mapped entries to be downloaded. Then the AD Agent initiates a UDP connection to the ASA and sends the change of authorization request packet.

  If for some reason the packet is lost, there is no way for the ASA to discern this. As a result, the ASA holds the session for 4-5 minutes, during which time this error message continues to appear if you have issued the **user-identity update active-user-database** command.
- When you use the Cisco Context Directory Agent (CDA) in conjunction with the ASA or Cisco Ironport Web Security Appliance (WSA), make sure that you open the following ports:
  - Authentication port for UDP—1645
  - Accounting port for UDP—1646

- Listening port for UDP—3799

  The listening port is used to send change of authorization requests from the CDA to the ASA or to the WSA.

- If the **user-identity action domain-controller-down** *domain_name* **disable user-identity-rule** command is configured and the specified domain is down, or if the **user-identity action ad-agent-down disable user-identity-rule** command is configured and the AD Agent is down, all the logged-in users have the disabled status.
- For domain names, the following characters are not valid: \/:*?"<>|.
- For usernames, the following characters are not valid: \/[]:;=,+*?"<>|@.
- For user group names, the following characters are not valid: \/[]:;=,+*?"<>|.
- How you configure the Identity Firewall to retrieve user information from the AD Agent affects the amount of memory used by the feature. You specify whether the ASA uses on-demand retrieval or full download retrieval. Choosing on-demand retrieval has the benefit of using less memory, because only users of received packets are queried and stored.

Prerequisites for the Identity Firewall

This section lists the prerequisites for configuring the Identity Firewall.

AD Agent

- The AD Agent must be installed on a Windows server that is accessible to the ASA. Additionally, you must configure the AD Agent to obtain information from the Active Directory servers and to communicate with the ASA.
- Supported Windows servers include Windows 2003, Windows 2008, and Windows 2008 R2.

Note Windows 2003 R2 is not supported for the AD Agent server.

- For the steps to install and configure the AD Agent, see the *Installation and Setup Guide for the Active Directory Agent*.
- Before configuring the AD Agent in the ASA, obtain the secret key value that the AD Agent and the ASA use to communicate. This value must match on both the AD Agent and the ASA.

Microsoft Active Directory

- Microsoft Active Directory must be installed on a Windows server and accessible by the ASA. Supported versions include Windows 2003, 2008, and 2008 R2 servers.
- Before configuring the Active Directory server on the ASA, create a user account in Active Directory for the ASA.

- Additionally, the ASA sends encrypted log-in information to the Active Directory server by using SSL enabled over LDAP. SSL must be enabled on the Active Directory server. See the documentation for Microsoft Active Directory for how to enable SSL for Active Directory.



Note Before running the AD Agent Installer, you must install the patches listed in the *README First for the Cisco Active Directory Agent* on each Microsoft Active Directory server that the AD Agent monitors. These patches are required even when the AD Agent is installed directly on the domain controller server.

Configure the Identity Firewall

To configure the Identity Firewall, perform the following tasks:

Procedure

Step 1 Configure the Active Directory domain in the ASA.

Step 2 Configure the AD Agent in ASA.

Step 3 Configure Identity Options.

Step 4 Configure Identity-based Security Policy. After the AD domain and AD Agent are configured, you can create identity-based object groups and ACLs for use in many features.

Configure the Active Directory Domain

Active Directory domain configuration on the ASA is required for the ASA to download Active Directory groups and accept user identities from specific domains when receiving IP-user mapping from the AD Agent.

Before you begin

- Active Directory server IP address
- Distinguished Name for LDAP base DN
- Distinguished Name and password for the Active Directory user that the Identity Firewall uses to connect to the Active Directory domain controller

To configure the Active Directory domain, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Identity Options**.

Step 2 Check the **Enable User Identity** check box to enable user identity.

Step 3 Click **Add**.

The **Domain** dialog box appears.

Step 4 Enter a domain name of up to 32 characters consisting of [a-z], [A-Z], [0-9], [!@#$%^&()-_=+[]{};,.] except '.' and ' ' at the first character. If the domain name includes a space, you must enclose that space character in quotation marks. The domain name is not case sensitive.

When you edit the name of an existing domain, the domain name associated with existing users and user groups is not changed.

Step 5 Select the Active Directory servers to associate with this domain, or click **Manage** to add a new server group to the list.

Step 6 Click **OK** to save the domain settings and close this dialog box.

Configure Active Directory Server Groups

To configure the Active Directory server group, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Identity Options** > **Add** > **Manage**.

The **Configure Active Directory Server Groups** dialog box appears.

Step 2 Click **Add**.

The **Add Active Directory Server Group** dialog box appears.

Step 3 To add servers to an Active Directory server group, select the group from the **Active Directory Server Groups** list, then click **Add**.

The **Add Active Directory Server** dialog box appears.

Step 4 Click **OK** to save the settings and close this dialog box.

Configure Active Directory Agents

Before you begin

- AD agent IP address
- Shared secret between the ASA and AD agent

To configure the AD Agents, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Identity Options**.

Step 2 Check the **Enable User Identity** check box to enable the feature.

Step 3 Click **Manage** in the **Active Directory Agent** section.

The **Configure Active Directory Agents** dialog box appears.

Step 4 Click the **Add** button.

Step 5 Click **OK** to save your changes and close this dialog box.

Configure Active Directory Agent Groups

Configure the primary and secondary AD Agents for the AD Agent Server Group. When the ASA detects that the primary AD Agent is not responding and a secondary agent is specified, the ASA switches to the secondary AD Agent. The Active Directory server for the AD agent uses RADIUS as the communication protocol; therefore, you should specify a key attribute for the shared secret between the ASA and AD Agent.

To configure the AD Agent Groups, perform the following steps:

Procedure

Step 1 Click **Add** from the **Configure Active Directory Agents** dialog box.

The **Add Active Directory Agent Group** dialog box appears.

Step 2 Enter a name for the AD Agent group.

Step 3 Specify the interface on which the ASA listens for traffic from the AD Agent server, and enter the FQDN of the server or IP address in the **Primary Active Directory Agent** section.

Step 4 Enter a timeout interval and the retry interval for the attempts that the ASA will continue to contact the AD Agent when it is not responding in the **Primary Active Directory Agent** section.

Step 5 Enter the shared secret key that is used between the primary AD Agent and the ASA.

Step 6 Specify the interface on which the ASA listens for traffic from the AD Agent server, and enter the FQDN of the server or IP address in the **Secondary Active Directory Agent** section.

Step 7 Enter a timeout interval and the retry interval for the attempts that the ASA will continue to perform to contact the AD Agent when it is not responding in the **Secondary Active Directory Agent** section.

Step 8 Enter the shared secret key that is used between the secondary AD Agent and the ASA.

Step 9 Click **OK** to save your changes and close this dialog box.

Configure Identity Options

To configure the Identity Options for the Identity Firewall, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Identity Options**.

Step 2 Check the **Enable User Identity** check box.

Step 3 To add a domain for the Identity Firewall, click **Add** to display the **Add Domain** dialog box.

Step 4 For domains that have already been added to the Domains list, check whether or not to disable rules when the domain is down because the Active Directory domain controller is not responding.

When a domain is down and this option is checked for that domain, the ASA disables the user identity rules associated with the users in that domain. Additionally, the status of all user IP addresses in that domain is marked as disabled in the **Monitoring** > **Properties** > **Identity** > **Users** pane.

Step 5 Choose the default domain for the Identity Firewall.

The default domain is used for all users and user groups when a domain has not been explicitly configured for those users or groups. When a default domain is not specified, the default domain for users and groups is LOCAL.

Additionally, the Identity Firewall uses the LOCAL domain for all locally defined user groups or locally defined users (those who log in and authenticate by using a VPN or web portal).

Note
The default domain name that you select must match the NetBIOS domain name configured on the Active Directory domain controller. If the domain name does not match, the AD Agent incorrectly associates the user-IP mapping with the domain name that you entered when configuring the ASA. To view the NetBIOS domain name, open the Active Directory user event security log in any text editor.

For multiple context modes, you can set a default domain name for each context, as well as within the system execution space.

Step 6 Choose the AD Agent group from the drop-down list. Click **Manage** to add AD Agent groups.

Step 7 Enter a number between 10 to 65535 seconds in the **Hello Timer** field.

The hello timer between the ASA and the AD Agent defines how frequently the ASA exchanges hello packets. The ASA uses the hello packet to obtain ASA replication status (in-sync or out-of-sync) and domain status (up or down). If the ASA does not receive a response from the AD Agent, it resends a hello packet after the specified interval.

Specify the number of times that the ASA will continue to send hello packets to the AD Agent. By default, the number of seconds is set to 30 and the retry times is set to 5.

Step 8 Check the **Enable Event Timestamp** check box to enable the ASA to keep track of the last event time stamp that it receives for each identifier and to discard any message if the event time stamp is at least 5 minutes older than the ASA's clock, or if its time stamp is earlier than the last event's time stamp.

For a newly booted ASA that does not have knowledge of the last event time stamp, the ASA compares the event time stamp with its own clock. If the event is at least 5 minutes older, the ASA does not accept the message.

We recommend that you configure the ASA, Active Directory, and Active Directory agent to synchronize their clocks among themselves using NTP

Step 9 Enter the number of hours in the **Poll Group Timer** field that the ASA uses to query the DNS server to resolve fully qualified domain names (FQDN). By default, the poll timer is set to 4 hours.

Step 10 Choose an option from the list in the **Retrieve User Information** section:

- **On Demand**—Specifies that the ASA retrieve the user mapping information of an IP address from the AD Agent when the ASA receives a packet that requires a new connection and the user of its source IP address is not in the user-identity database.
- **Full Download**—Specifies that the ASA send a request to the AD Agent to download the entire IP-user mapping table when the ASA starts and then to receive incremental IP-user mapping when users log in and log out.

  Note
  Choosing **On Demand** has the benefit of using less memory because only users of received packets are queried and stored.

Step 11 Choose whether or not to disable rules if the AD Agent is not responding.

When the AD Agent is down and this option is selected, the ASA disables the user identity rules associated with the users in that domain. Additionally, the status of all user IP addresses in that domain are marked as disabled in the **Monitoring** > **Properties** > **Identity** > **Users** pane.

Step 12 Choose whether or not to remove a user's IP address when the NetBIOS probe fails.

Choosing this option specifies the action when NetBIOS probing to a user is blocked (for example, the user client does not respond to a NetBIOS probe). The network connection might be blocked to that client or the client is not active. When this option is chosen, the ASA disables the identity rules associated with that user's IP address.

Step 13 Choose whether or not to remove a user's MAC address when it is inconsistent with the IP address that the ASA has currently mapped to that MAC address. When this option is chosen, the ASA disables the user identity rules associated with the specific user.

Step 14 Choose whether or not to track users that are not found.

Step 15 Choose the **Idle Timeout** option and enter a time in minutes, from 1 minute to 65535. By default, the idle timeout is set to 60 minutes.

Enabling this option configures a timer when an active user is considered idle, meaning the ASA does not receive traffic from the user's IP address for more than the specified time. After the timer expires, the user's IP address is marked inactive and removed from the local cached IP-user database and the ASA no longer notifies the AD Agent about that IP address. Existing traffic is still allowed to pass. When the **Idle Timeout** option is enabled, the ASA runs an inactive timer even when the NetBIOS Logout Probe is configured.

Note
The **Idle Timeout** option does not apply to VPN or cut-through proxy users.

Step 16 Enable NetBIOS probing and set the probe timer (from 1 to 65535 minutes) before a user's IP addresses is probed and the retry interval (from 1 to 256 retries) between retry probes.

Enabling this option configures how often the ASA probes the user host to determine whether the user client is still active. To minimize the NetBIOS packets, ASA only sends a NetBIOS probe to the client when the user has been idle for more than the specified number of minutes in the Idle Timeout minutes field.

Step 17 Choose an option from the **User Name** list:

- **Match Any**—As long as the NetBIOS response from the host includes the username of the user assigned to the IP address, the user identity is be considered valid. Specifying this option requires that the host enabled the Messenger service and configured a WINS server.
- **Exact Match**—The username of the user assigned to the IP address must be the only one in the NetBIOS response. Otherwise, the user identity of that IP address is considered invalid. Specifying this option requires that the host enabled the Messenger service and configured a WINS server.
- **User Not Needed**—As long as the ASA received a NetBIOS response from the host, the user identity is considered valid.

Step 18 Click **Apply** to save the Identity Firewall configuration.

Configure Identity-Based Security Policy

You can incorporate identity-based policy in many ASA features. Any feature that uses extended ACLs (other than those listed as unsupported in the Guidelines section) can take advantage of an identity firewall. You can now add user identity arguments to extended ACLs, as well as network-based parameters.

Features that can use identity include the following:

- Access rules—An access rule permits or denies traffic on an interface using network information. With an identity firewall, you can control access based on user identity.
- AAA rules—An authentication rule (also known as cut-through proxy) controls network access based on the user. Because this function is very similar to an access rule plus an identity firewall, AAA rules can now be used as a backup method of authentication if a user's AD login expires. For example, for any user without a valid login, you can trigger a AAA rule. To ensure that the AAA rule is only triggered for users that do not have valid logins, you can specify special usernames in the extended ACL used for the access rule and for the AAA rule: None (users without a valid login) and Any (users with a valid login). In the access rule, configure your policy as usual for users and groups, but then include a AAA rule that permits all None users; you must permit these users so they can later trigger a AAA rule. Then, configure a AAA rule that denies Any users (these users are not subject to the AAA rule, and were handled already by the access rule), but permits all None users. For example:

```
access-list 100 ex permit ip user CISCO\xyz any any
access-list 100 ex deny ip user CISCO\abc any any
access-list 100 ex permit ip user NONE any any
access-list 100 ex deny any any
access-group 100 in interface inside

access-list 200 ex deny ip user ANY any any
access-list 200 ex permit user NONE any any
aaa authenticate match 200 inside user-identity
```

  For more information, see the legacy feature guide.

- VPN filter—Although a VPN does not support identity firewall ACLs in general, you can configure the ASA to enforce identity-based access rules on VPN traffic. By default, VPN traffic is not subject to access rules. You can force VPN clients to abide by access rules that use an identity firewall ACL (with the **no sysopt connection permit-vpn** command). You can also use an identity firewall ACL with the VPN filter feature; a VPN filter accomplishes a similar effect by allowing access rules in general.

Monitoring the Identity Firewall

See the following screens for monitoring the Identity Firewall status:

- **Monitoring** > **Properties** > **Identity** > **AD Agent**

  This pane shows the status of the AD Agents and the domains, and the statistics for the AD Agents.

- **Monitoring** > **Properties** > **Identity** > **Memory Usage**

  This pane shows the memory usage that the Identity Firewall consumes on the ASA.

- **Monitoring** > **Properties** > **Identity** > **User**

- This pane shows information about all users contained in the IP-user mapping database used by the Identity Firewall.

- **Monitoring** > **Properties** > **Identity** > **Group**

  This pane shows the list of user groups configured for the Identity Firewall.

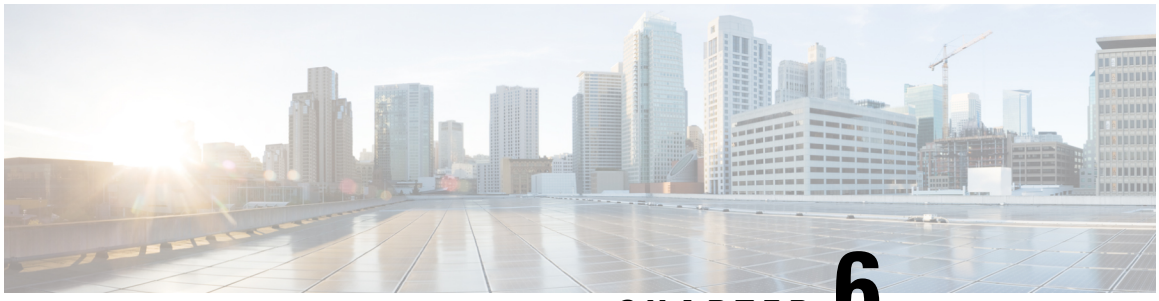
- **Tools** > **Command Line Interface**

  This pane allows you to issue various non-interactive commands and view results.

History for the Identity Firewall

Table 3: History for the Identity Firewall

Feature Name	Releases	Description
Identity Firewall	8.4(2)	The Identity Firewall feature was introduced. We introduced or modified the following screens: **Configuration** > **Firewall** > **Identity Options** **Configuration** > **Firewall** > **Objects** > **Local User Groups** **Monitoring** > **Properties** > **Identity**.



CHAPTER 6

ASA and Cisco TrustSec

This chapter describes how to implement Cisco TrustSec for the ASA.

- About Cisco TrustSec, on page 73
- Guidelines for Cisco TrustSec, on page 81
- Configure the ASA to Integrate with Cisco Trustsec, on page 83
- AnyConnect Client VPN Support for Cisco TrustSec, on page 92
- Monitoring Cisco TrustSec, on page 94
- History for Cisco TrustSec, on page 95

About Cisco TrustSec

Traditionally, security features such as firewalls performed access control based on predefined IP addresses, subnets, and protocols. However, with enterprises transitioning to borderless networks, both the technology used to connect people and organizations and the security requirements for protecting data and networks have evolved significantly. Endpoints are becoming increasingly nomadic and users often employ a variety of endpoints (for example, laptop versus desktop, smart phone, or tablet), which means that a combination of user attributes plus endpoint attributes provide the key characteristics (in addition to existing 6-tuple based rules), that enforcement devices such as switches and routers with firewall features or dedicated firewalls can reliably use for making access control decisions.

As a result, the availability and propagation of endpoint attributes or client identity attributes have become increasingly important requirements to enable security across the customers' networks, at the access, distribution, and core layers of the network, and in the data center.

Cisco TrustSec provides access control that builds upon an existing identity-aware infrastructure to ensure data confidentiality between network devices and integrate security access services on one platform. In the Cisco TrustSec feature, enforcement devices use a combination of user attributes and endpoint attributes to make role-based and identity-based access control decisions. The availability and propagation of this information enables security across networks at the access, distribution, and core layers of the network.

Implementing Cisco TrustSec into your environment has the following advantages:

- Provides a growing mobile and complex workforce with appropriate and more secure access from any device
- Lowers security risks by providing comprehensive visibility of who and what is connecting to the wired or wireless network
- Offers exceptional control over activity of network users accessing physical or cloud-based IT resources

- Reduces total cost of ownership through centralized, highly secure access policy management and scalable enforcement mechanisms
- For more information, see the following URLs:
  - Description of the Cisco TrustSec system and architecture for the enterprise.

    http://www.cisco.com/c/en/us/solutions/enterprise-networks/trustsec/index.html
  - Instructions for deploying the Cisco TrustSec solution in the enterprise, including links to component design guides.

    http://www.cisco.com/c/en/us/solutions/enterprise/design-zone-security/landing_DesignZone_TrustSec.html
  - An overview of the Cisco TrustSec solution when used with the ASA, switches, wireless LAN (WLAN) controllers, and routers.

    http://www.cisco.com/c/en/us/solutions/collateral/enterprise-networks/trustsec/solution_overview_c22-591771.pdf
  - The Cisco TrustSec Platform Support Matrix, which lists the Cisco products that support the Cisco TrustSec solution.

    http://www.cisco.com/c/en/us/solutions/enterprise-networks/trustsec/trustsec_matrix.html

About SGT and SXP Support in Cisco TrustSec

In the Cisco TrustSec feature, security group access transforms a topology-aware network into a role-based network, which enables end-to-end policies enforced on the basis of role-based access control (RBAC). Device and user credentials acquired during authentication are used to classify packets by security groups. Every packet entering the Cisco TrustSec cloud is tagged with a security group tag (SGT). The tagging helps trusted intermediaries identify the source identity of the packet and enforce security policies along the data path. An SGT can indicate a privilege level across the domain when the SGT is used to define a security group ACL.

An SGT is assigned to a device through IEEE 802.1X authentication, web authentication, or MAC authentication bypass (MAB), which occurs with a RADIUS vendor-specific attribute. An SGT can be assigned statically to a particular IP address or to a switch interface. An SGT is passed along dynamically to a switch or access point after successful authentication.

The Security-group eXchange Protocol (SXP) is a protocol developed for Cisco TrustSec to propagate the IP-to-SGT mapping database across network devices that do not have SGT-capable hardware support to hardware that supports SGTs and security group ACLs. SXP, a control plane protocol, passes IP-SGT mapping from authentication points (such as legacy access layer switches) to upstream devices in the network.

The SXP connections are point-to-point and use TCP as the underlying transport protocol. SXP uses TCP port 64999 to initiate a connection. Additionally, an SXP connection is uniquely identified by the source and destination IP addresses.

Roles in the Cisco TrustSec Feature

To provide identity and policy-based access enforcement, the Cisco TrustSec feature includes the following roles:

- Access Requester (AR)—Access requesters are endpoint devices that request access to protected resources in the network. They are primary subjects of the architecture and their access privilege depends on their Identity credentials.

Access requesters include endpoint devices such PCs, laptops, mobile phones, printers, cameras, and MACsec-capable IP phones.

- Policy Decision Point (PDP)—A policy decision point is responsible for making access control decisions. The PDP provides features such as 802.1x, MAB, and web authentication. The PDP supports authorization and enforcement through VLAN, DACL, and security group access (SGACL/SXP/SGT).

  In the Cisco TrustSec feature, the Cisco Identity Services Engine (ISE) acts as the PDP. The Cisco ISE provides identity and access control policy functionality.

- Policy Information Point (PIP)—A policy information point is a source that provides external information (for example, reputation, location, and LDAP attributes) to policy decision points.

  Policy information points include devices such as Session Directory, Sensor IPS, and Communication Manager.

- Policy Administration Point (PAP)—A policy administration point defines and inserts policies into the authorization system. The PAP acts as an identity repository by providing Cisco TrustSec tag-to-user identity mapping and Cisco TrustSec tag-to-server resource mapping.

  In the Cisco TrustSec feature, the Cisco Secure Access Control System (a policy server with integrated 802.1x and SGT support) acts as the PAP.

- Policy Enforcement Point (PEP)—A policy enforcement point is the entity that carries out the decisions (policy rules and actions) made by the PDP for each AR. PEP devices learn identity information through the primary communication path that exists across networks. PEP devices learn the identity attributes of each AR from many sources, such as endpoint agents, authorization servers, peer enforcement devices, and network flows. In turn, PEP devices use SXP to propagate IP-SGT mapping to mutually trusted peer devices across the network.

  Policy enforcement points include network devices such as Catalyst switches, routers, firewalls (specifically the ASA), servers, VPN devices, and SAN devices.

The ASA serves the PEP role in the identity architecture. Using SXP, the ASA learns identity information directly from authentication points and uses it to enforce identity-based policies.

Security Group Policy Enforcement

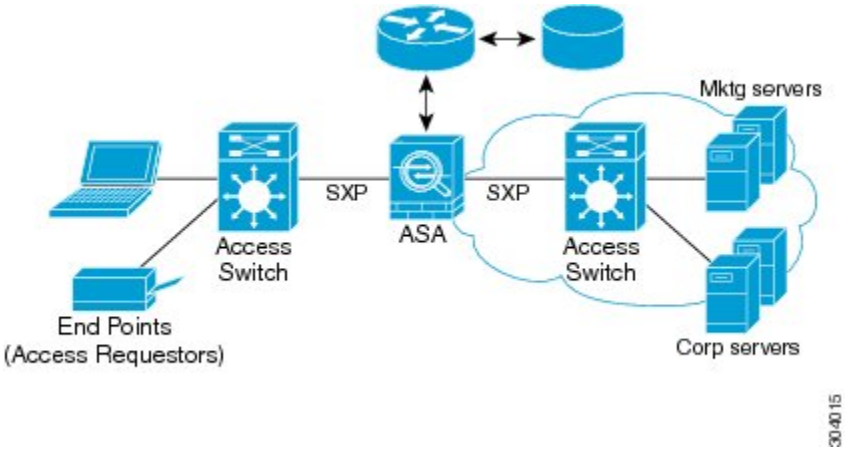
Security policy enforcement is based on security group name. An endpoint device attempts to access a resource in the data center. Compared to traditional IP-based policies configured on firewalls, identity-based policies are configured based on user and device identities. For example, mktg-contractor is allowed to access mktg-servers; mktg-corp-users are allowed to access mktg-server and corp-servers.

The benefits of this type of deployment include the following:

- User group and resource are defined and enforced using single object (SGT) simplified policy management.
- User identity and resource identity are retained throughout the Cisco TrustSec-capable switch infrastructure.

The following figure shows a deployment for security group name-based policy enforcement.

Figure 10: Security Group Name-Based Policy Enforcement Deployment



Implementing Cisco TrustSec allows you to configure security policies that support server segmentation and includes the following features:

- A pool of servers can be assigned an SGT for simplified policy management.
- The SGT information is retained within the infrastructure of Cisco TrustSec-capable switches.
- The ASA can use the IP-SGT mapping for policy enforcement across the Cisco TrustSec domain.
- Deployment simplification is possible because 802.1x authorization for servers is mandatory.

How the ASA Enforces Security Group-Based Policies

Note User-based security policies and security-group based policies can coexist on the ASA. Any combination of network, user-based, and security-group based attributes can be configured in a security policy.

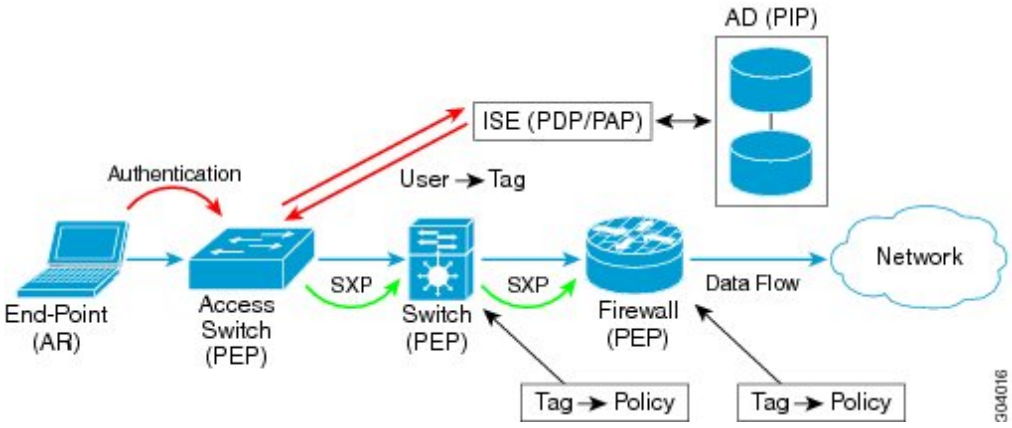
To configure the ASA to function with Cisco TrustSec, you must import a Protected Access Credential (PAC) file from the ISE.

Importing the PAC file to the ASA establishes a secure communication channel with the ISE. After the channel is established, the ASA initiates a PAC secure RADIUS transaction with the ISE and downloads Cisco TrustSec environment data (that is, the security group table). The security group table maps SGTs to security group names. Security group names are created on the ISE and provide user-friendly names for security groups.

The first time that the ASA downloads the security group table, it walks through all entries in the table and resolves all the security group names included in security policies that have been configured on it; then the ASA activates those security policies locally. If the ASA cannot resolve a security group name, it generates a syslog message for the unknown security group name.

The following figure shows how a security policy is enforced in Cisco TrustSec.

Figure 11: Security Policy Enforcement



1. An endpoint device connects to an access layer device directly or via remote access and authenticates with Cisco TrustSec.
2. The access layer device authenticates the endpoint device with the ISE by using authentication methods such as 802.1X or web authentication. The endpoint device passes role and group membership information to classify the device into the appropriate security group.
3. The access layer device uses SXP to propagate the IP-SGT mapping to the upstream devices.
4. The ASA receives the packet and looks up the SGTs for the source and destination IP addresses using the IP-SGT mapping passed by SXP.

   If the mapping is new, the ASA records it in its local IP-SGT Manager database. The IP-SGT Manager database, which runs in the control plane, tracks IP-SGT mapping for each IPv4 or IPv6 address. The database records the source from which the mapping was learned. The peer IP address of the SXP connection is used as the source of the mapping. Multiple sources can exist for each IP-SGT mapped entry.

   If the ASA is configured as a Speaker, the ASA transmits all IP-SGT mapping entries to its SXP peers.
5. If a security policy is configured on the ASA with that SGT or security group name, the ASA enforces the policy. (You can create security policies on the ASA that include SGTs or security group names. To enforce policies based on security group names, the ASA needs the security group table to map security group names to SGTs.)

   If the ASA cannot find a security group name in the security group table and it is included in a security policy, the ASA considers the security group name to be unknown and generates a syslog message. After the ASA refreshes the security group table from the ISE and learns the security group name, the ASA generates a syslog message indicating that the security group name is known.

Effects of Changes to Security Groups on the ISE

The ASA periodically refreshes the security group table by downloading an updated table from the ISE. Security groups can change on the ISE between downloads. These changes are not reflected on the ASA until it refreshes the security group table.

Tip We recommend that you schedule policy configuration changes on the ISE during a maintenance window, then manually refresh the security group table on the ASA to make sure the security group changes have been incorporated.

Handling policy configuration changes in this way maximizes the chances of security group name resolution and immediate activation of security policies.

The security group table is automatically refreshed when the environment data timer expires. You can also trigger a security group table refresh on demand.

If a security group changes on the ISE, the following events occur when the ASA refreshes the security group table:

- Only security group policies that have been configured using security group names need to be resolved with the security group table. Policies that include security group tags are always active.
- When the security group table is available for the first time, all policies with security group names are walked through, security group names are resolved, and policies are activated. All policies with tags are walked through, and syslogs are generated for unknown tags.
- If the security group table has expired, policies continue to be enforced according to the most recently downloaded security group table until you clear it, or a new table becomes available.
- When a resolved security group name becomes unknown on the ASA, it deactivates the security policy; however, the security policy persists in the ASA running configuration.
- If an existing security group is deleted on the PAP, a previously known security group tag can become unknown, but no change in policy status occurs on the ASA. A previously known security group name can become unresolved, and the policy is then inactivated. If the security group name is reused, the policy is recompiled using the new tag.
- If a new security group is added on the PAP, a previously unknown security group tag can become known, a syslog message is generated, but no change in policy status occurs. A previously unknown security group name can become resolved, and associated policies are then activated.
- If a tag has been renamed on the PAP, policies that were configured using tags display the new name, and no change in policy status occurs. Policies that were configured with security group names are recompiled using the new tag value.

Speaker and Listener Roles on the ASA

The ASA supports SXP to send and receive IP-SGT mapping entries to and from other network devices. Using SXP allows security devices and firewalls to learn identity information from access switches without the need for hardware upgrades or changes. SXP can also be used to pass IP-SGT mapping entries from upstream devices (such as data center devices) back to downstream devices. The ASA can receive information from both upstream and downstream directions.

When configuring an SXP connection on the ASA to an SXP peer, you must designate the ASA as a Speaker or a Listener for that connection so that it can exchange Identity information:

- Speaker mode—Configures the ASA so that it can forward all active IP-SGT mapping entries collected on the ASA to upstream devices for policy enforcement.

- Listener mode—Configures the ASA so that it can receive IP-SGT mapping entries from downstream devices (SGT-capable switches) and use that information to create policy definitions.

If one end of an SXP connection is configured as a Speaker, then the other end must be configured as a Listener, and vice versa. If both devices on each end of an SXP connection are configured with the same role (either both as Speakers or both as Listeners), the SXP connection fails and the ASA generates a syslog message.

Multiple SXP connections can learn IP-SGT mapping entries that have been downloaded from the IP-SGT mapping database. After an SXP connection to an SXP peer is established on the ASA, the Listener downloads the entire IP-SGT mapping database from the Speaker. All changes that occur after this are sent only when a new device appears on the network. As a result, the rate of SXP information flow is proportional to the rate at which end hosts authenticate to the network.

IP-SGT mapping entries that have been learned through SXP connections are maintained in the SXP IP-SGT mapping database. The same mapping entries may be learned through different SXP connections. The mapping database maintains one copy for each mapping entry learned. Multiple mapping entries of the same IP-SGT mapping value are identified by the peer IP address of the connection from which the mapping was learned. SXP requests that the IP-SGT Manager add a mapping entry when a new mapping is learned the first time and remove a mapping entry when the last copy in the SXP database is removed.

Whenever an SXP connection is configured as a Speaker, SXP requests that the IP-SGT Manager forward all the mapping entries collected on the device to the peer. When a new mapping is learned locally, the IP-SGT Manager requests that SXP forward it through connections that are configured as Speakers.

Configuring the ASA to be both a Speaker and a Listener for an SXP connection can cause SXP looping, which means that SXP data can be received by an SXP peer that originally transmitted it.

Register the ASA with the ISE

The ASA must be configured as a recognized Cisco TrustSec network device in the ISE before the ASA can successfully import a PAC file. To register the ASA with the ISE, perform the following steps:

Procedure

Step 1 Log into the ISE.

Step 2 Choose **Administration** > **Network Devices** > **Network Devices**.

Step 3 Click **Add**.

Step 4 Enter the IP address of the ASA.

Step 5 When the ISE is being used for user authentication, enter a shared secret in the Authentication Settings area.

When you configure the AAA sever on the ASA, provide the shared secret that you create here on the ISE. The AAA server on the ASA uses this shared secret to communicate with the ISE.

Step 6 Specify a device name, device ID, password, and a download interval for the ASA. See the ISE documentation for how to perform these tasks.

Create a Security Group on the ISE

When configuring the ASA to communicate with the ISE, you specify a AAA server. When configuring the AAA server on the ASA, you must specify a server group. The security group must be configured to use the RADIUS protocol. To create a security group on the ISE, perform the following steps:

Procedure

Step 1 Log into the ISE.

Step 2 Choose **Policy** > **Policy Elements** > **Results** > **Security Group Access** > **Security Group**.

Step 3 Add a security group for the ASA. (Security groups are global and not ASA specific.)

The ISE creates an entry under Security Groups with a tag.

Step 4 In the Security Group Access area, configure device ID credentials and a password for the ASA.

Generate the PAC File

To generate the PAC file, perform the following steps.



Note The PAC file includes a shared key that allows the ASA and ISE to secure the RADIUS transactions that occur between them. For this reason, make sure that you store it securely on the ASA.

Procedure

Step 1 Log into the ISE.

Step 2 Choose **Administration** > **Network Resources** > **Network Devices**.

Step 3 From the list of devices, choose the ASA.

Step 4 Under the Security Group Access (SGA), click **Generate PAC**.

Step 5 To encrypt the PAC file, enter a password.

The password (or encryption key) that you enter to encrypt the PAC file is independent of the password that was configured on the ISE as part of the device credentials.

The ISE generates the PAC file. The ASA can import the PAC file from flash or from a remote server via TFTP, FTP, HTTP, HTTPS, or SMB. (The PAC file does not have to reside on the ASA flash before you can import it.)

Guidelines for Cisco TrustSec

This section includes the guidelines and limitations that you should review before configuring Cisco TrustSec.

Failover

- You can configure security group-based policies on the ASA in both the Active/Active and Active/Standby configurations.
- When the ASA is part of a failover configuration, you must import the PAC file to the primary ASA device. You must also refresh the environment data on the primary device.
- The ASA can communicate with the ISE configured for high availability (HA).
- You can configure multiple ISE servers on the ASA and if the first server is unreachable, it continues to the next server, and so on. However, if the server list is downloaded as part of the Cisco TrustSec environment data, it is ignored.
- If the PAC file downloaded from the ISE expires on the ASA and it cannot download an updated security group table, the ASA continues to enforce security policies based on the last downloaded security group table until the ASA downloads an updated table.

Clustering

- When the ASA is part of a clustering configuration, you must import the PAC file to the control unit.
- When the ASA is part of a clustering configuration, you must refresh the environment data on the control unit.

IPv6

The ASA supports SXP for IPv6 and IPv6-capable network devices. The AAA server must use an IPv4 address.

Layer 2 SGT Imposition

- Supported only on physical interfaces, subinterfaces, redundant interfaces, and EtherChannel interfaces.
- Not supported on logical interfaces or virtual interfaces, such as a BVI.
- Does not support link encryption using SAP negotiation and MACsec.
- Not supported on failover links.
- Not supported on cluster control links.
- The ASA does not reclassify existing flows if the SGT is changed. Any policy decisions that were made based on the previous SGT remain in force for the life of the flow. However, the ASA can immediately reflect SGT changes on egress packets, even if the packets belong to a flow whose classification was based on a previous SGT.
- Firepower 1010 switch ports and VLAN interfaces do not support Layer 2 Security Group Tagging Imposition.

Additional Guidelines

- The ASA supports SXP Version 3. The ASA negotiates SXP versions with different SXP-capable network devices.
- The ASA does not support Change of Authorization (CoA) for the Cisco TrustSec (CTS) environment data.
- You can configure the ASA to refresh the security group table when the SXP reconcile timer expires and you can download the security group table on demand. When the security group table on the ASA is updated from the ISE, changes are reflected in the appropriate security policies.
- Cisco TrustSec supports the Smart Call Home feature in single context and multi-context mode, but not in the system context.
- The ASA can only be configured to interoperate in a single Cisco TrustSec domain.
- The ASA does not support static configuration of SGT-name mapping on the device.
- NAT is not supported in SXP messages.
- SXP conveys IP-SGT mapping to enforcement points in the network. If an access layer switch belongs to a different NAT domain than the enforcing point, the IP-SGT map that it uploads is invalid, and an IP-SGT mapping database lookup on the enforcement device does not yield valid results. As a result, the ASA cannot apply security group-aware security policy on the enforcement device.
- You can configure a default password for the ASA to use for SXP connections, or you can choose not to use a password; however, connection-specific passwords are not supported for SXP peers. The configured default SXP password should be consistent across the deployment network. If you configure a connection-specific password, connections may fail and a warning message appears. If you configure the connection with the default password, but it is not configured, the result is the same as when you have configured the connection with no password.
- The ASA can be configured as an SXP Speaker or Listener, or both. However, SXP connection loops can form when a device has bidirectional connections to a peer or is part of a unidirectionally connected chain of devices. (The ASA can learn IP-SGT mapping for resources from the access layer in the data center. The ASA might need to propagate these tags to downstream devices.) SXP connection loops can cause unexpected behavior of SXP message transport. In cases where the ASA is configured to be a Speaker and Listener, an SXP connection loop can occur, causing SXP data to be received by the peer that originally transmitted it.
- When changing the ASA local IP address, you must ensure that all SXP peers have updated their peer list. In addition, if SXP peers changes its IP addresses, you must ensure those changes are reflected on the ASA.
- Automatic PAC file provisioning is not supported. The ASA administrator must request the PAC file from the ISE administrative interface and import it into the ASA.
- PAC files have expiration dates. You must import the updated PAC file before the current PAC file expires; otherwise, the ASA cannot retrieve environment data updates. If the PAC file downloaded from the ISE expires on the ASA and it cannot download an updated security group table, the ASA continues to enforce security policies based on the last downloaded security group table until the ASA downloads an updated table.
- When a security group changes on the ISE (for example, it is renamed or deleted), the ASA does not change the status of any ASA security policies that contain an SGT or security group name associated

with the changed security group; however, the ASA generates a syslog message to indicate that those security policies changed.

- The multi-cast types are not supported in ISE 1.0.
- An SXP connection stays in the initializing state among two SXP peers interconnected by the ASA; as shown in the following example:

```
(SXP peer A) - - - - (ASA) - - - (SXP peer B)
```

Therefore, when configuring the ASA to integrate with Cisco TrustSec, you must enable the no-NAT, no-SEQ-RAND, and MD5-AUTHENTICATION TCP options on the ASA to configure SXP connections. Create a TCP state bypass policy for traffic destined to SXP port TCP 64999 among the SXP peers. Then apply the policy on the appropriate interfaces.

For example, the following set of commands shows how to configure the ASA for a TCP state bypass policy:

```
access-list SXP-MD5-ACL extended permit tcp host peerA host peerB eq 64999
access-list SXP-MD5-ACL extended permit tcp host peerB host peerA eq 64999

tcp-map SXP-MD5-OPTION-ALLOW
 tcp-options range 19 19 allow

class-map SXP-MD5-CLASSMAP
 match access-list SXP-MD5-ACL

policy-map type inspect dns preset_dns_map
 parameters
  message-length maximum 512
policy-map global_policy
class SXP-MD5-CLASSMAP
  set connection random-sequence-number disable
  set connection advanced-options SXP-MD5-OPTION-ALLOW
  set connection advanced-options tcp-state-bypass
service-policy global_policy global
```

Configure the ASA to Integrate with Cisco Trustsec

To configure the ASA to integrate with Cisco TrustSec, perform the following tasks.

Before you begin

Before configuring the ASA to integrate with Cisco TrustSec, you must complete the following tasks in ISE:

- Register the ASA with the ISE, on page 79
- Create a Security Group on the ISE, on page 80
- Generate the PAC File, on page 80

Procedure

Step 1 Configure the AAA Server for Cisco TrustSec Integration, on page 84

Step 2 Import a PAC File, on page 85

Step 3 Configure the Security Exchange Protocol, on page 86

This task enables and sets the default values for SXP.

Step 4 Add an SXP Connection Peer, on page 87

Step 5 Refresh Environment Data, on page 88

Do this as needed.

Step 6 Configure the Security Policy, on page 89

Step 7 Configure Layer 2 Security Group Tagging Imposition, on page 89

Configure the AAA Server for Cisco TrustSec Integration

This section describes how to integrate the AAA server for Cisco TrustSec. To configure the AAA server group to communicate with the ISE on the ASA, perform the following steps.

Before you begin

- The referenced server group must be configured to use the RADIUS protocol. If you add a non-RADIUS server group to the ASA, the configuration fails.
- If the ISE is also used for user authentication, obtain the shared secret that was entered on the ISE when you registered the ASA with the ISE. Contact your ISE administrator to obtain this information.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Identity By TrustSec**.

Step 2 Click **Manage** to add a server group to the ASA.

The **Configure AAA Server Group** dialog box appears.

Step 3 Enter the name of the security group that was created on the ISE for the ASA.

The server group name you specify here must match the name of the security group that was created on the ISE for the ASA. If these two group names do not match, the ASA cannot communicate with the ISE. Contact your ISE administrator to obtain this information.

Step 4 Choose **RADIUS** from the **Protocol** drop-down list.

To complete the remaining fields in the **AAA Server Group** dialog box, see the RADIUS chapter in the general operations configuration guide.

Step 5 Click **OK**.

Step 6 Select the AAA sever group that you just created and click **Add** in the **Servers in the Selected Group** area to add a server to a group.

The **Add AAA Server** dialog box appears.

Step 7 Select the network interface where the ISE server resides.

Step 8 Enter the IP address of the ISE server.

To complete the remaining fields in the AAA Server dialog box, see the RADIUS chapter in the general operations configuration guide.

Step 9 Click **OK**.

Step 10 Click **Apply** to save the changes to the running configuration.

Import a PAC File

This section describes how to import a PAC file.

Before you begin

- The ASA must be configured as a recognized Cisco TrustSec network device in the ISE before the ASA can generate a PAC file.
- Obtain the password used to encrypt the PAC file when generating it on the ISE. The ASA requires this password to import and decrypt the PAC file.
- When imported, the PAC file resides in NVRAM. When operating in HA mode, if you configure the failover and stateful links correctly, importing the PAC file into the active unit will result in replication to the secondary. Because the imported file resides in NVRAM, you must import the file again whenever the device reboots, for example, after a software upgrade.
- The device uses a single PAC file. If you import more than one, each imported PAC file replaces the previously imported file.
- The ASA requires access to the PAC file generated by the ISE. The ASA can import the PAC file from flash or from a remote server via TFTP, FTP, HTTP, HTTPS, or SMB. (The PAC file does not need to reside on the ASA flash before you can import it.)
- The server group has been configured for the ASA.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Identity By TrustSec**.

Step 2 Check the **Enable Security Exchange Protocol** check box to enable SXP.

Step 3 Click **Import PAC**to display the**Import PAC** dialog box.

Step 4 Enter the path and filename for the PAC file by using one of the following formats:

- **disk0**: Path and filename on disk0
- **disk1**: Path and filename on disk1

- **flash**: Path and filename on flash
- **ftp**: Path and filename on FTP
- **http**: Path and filename on HTTP
- **https**: Path and filename on HTTPS
- **smb**: Path and filename on SMB
- **tftp**: Path and filename on TFTP

Multi-mode

- **http**: Path and filename on HTTP
- **https**: Path and filename on HTTPS
- **smb**: Path and filename on SMB
- **tftp**: Path and filename on TFTP

Step 5 Enter the password used to encrypt the PAC file. The password is independent of the password that was configured on the ISE as part of the device credentials.

Step 6 Reenter the password to confirm it.

Step 7 Click **Import**.

Step 8 Click **Apply** to save the changes to the running configuration.

When you import the PAC file, the file is converted to ASCII HEX format and sent to the ASA in non-interactive mode.

Configure the Security Exchange Protocol

You need to enable and configure the Security Exchange Protocol (SXP) to use Cisco Trustsec.

Before you begin

At least one interface must be in the UP/UP state. If you enable SXP with all interfaces down, the ASA does not display a message indicating that SXP is not working or it could not be enabled. If you check the configuration by entering the **show running-config** command, the command output displays the following message:

```
"WARNING: SXP configuration in process, please wait for a few moments and try again."
```

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Identity By TrustSec**.

Step 2 Check the **Enable Security Exchange Protocol** check box to enable SXP. By default, SXP is disabled.

Step 3 (Optional; not recommended.) Enter the default local IP address for SXP connections. The IP address can be an IPv4 or IPv6 address.

Note
The ASA determines the local IP address for an SXP connection as the outgoing interface IP address that is reachable by the peer IP address. If the configured local address is different from the outgoing interface IP address, the ASA cannot connect to the SXP peer and generates a syslog message. We recommend that you do not configure a default source IP address for SXP connections and allow the ASA to perform a route/ARP lookup to determine the source IP address for an SXP connection.

Step 4 (Optional.) Enter the default password for TCP MD5 authentication with SXP peers. By default, SXP connections do not have a password set.

Configure a default password if and only if you configure the SXP connection peers to use the default password. The password can be up to 80 characters. It is not encrypted.

Step 5 (Optional.) Change the time interval between ASA attempts to set up new SXP connections between SXP peers in the **Retry Timer** field.

The ASA continues to make connection attempts until a successful connection is made, waiting the retry interval before trying again after a failed attempt. You can specify a retry period from 0 to 64000 seconds. The default is 120 seconds. If you specify 0 seconds, the ASA does not try to connect to SXP peers.

We recommend that you configure the retry timer to a different value from its SXP peer devices.

Step 6 (Optional.) Change the reconcile timer value.

After an SXP peer terminates its SXP connection, the ASA starts a hold down timer. If an SXP peer connects while the hold down timer is running, the ASA starts the reconciliation timer; then, the ASA updates the SXP mapping database to learn the latest mappings.

When the reconciliation timer expires, the ASA scans the SXP mapping database to identify stale mapping entries (which were learned in a previous connection session). The ASA marks these connections as obsolete. When the reconciliation timer expires, the ASA removes the obsolete entries from the SXP mapping database.

You can specify a reconciliation period from 1 to 64000 seconds. The default is 120 seconds.

Step 7 (Optional.) In **Network Map**, configure the depth of IPv4 subnet expansion when acting as a speaker to peers that use SXPv2 or lower.

If a peer uses SXPv2 or lower, the peer cannot understand SGT to subnet bindings. The ASA can expand the IPv4 subnet bindings to individual host bindings (IPv6 bindings are not expanded). This command specifies the maximum number of host bindings that can be generated from a subnet binding.

You can specify the maximum number to be from 0 to 65535. The default is 0, which means that subnet bindings are not expanded to host bindings.

Step 8 Click **Apply** to save the changes to the running configuration.

Add an SXP Connection Peer

To add an SXP connection peer, perform the following steps:

Procedure

Step 1 Choose **Configuration > Firewall** > **Identity By TrustSec**.

Step 2 Click **Add**to display the **Add Connection** dialog box.

Step 3 Enter the IPv4 or IPv6 address of the SXP peer. The peer IP address must be reachable from the ASA outgoing interface.

Step 4 Indicate whether or not to use the authentication key for the SXP connection by choosing one of the following values:

- **Default**—Use the default password configured for SXP connections.
- **None**—Do not use a password for the SXP connection.

Step 5 (Optional) Specify the mode of the SXP connection by choosing one of the following values:

- **Local**—Use the local SXP device.
- **Peer**—Use the peer SXP device.

Step 6 Specify whether the ASA functions as a Speaker or Listener for the SXP connection:

- **Speaker**—The ASA can forward IP-SGT mapping to upstream devices.
- **Listener**—The ASA can receive IP-SGT mapping from downstream devices.

Step 7 (Optional) Click **Advanced** and enter the local IPv4 or IPv6 address of the SXP connection.

The ASA uses a route lookup to determine the right interface. If you specify an address, it must match the route lookup interface address of the outbound interface. We recommend that you do not configure a source IP address for an SXP connection and allow the ASA to perform a route/ARP lookup to determine the source IP address for the SXP connection.

Step 8 Click **OK**.

Step 9 Click **Apply** to save your settings to the running configuration.

Refresh Environment Data

The ASA downloads environment data from the ISE, which includes the Security Group Tag (SGT) name table. The ASA automatically refreshes its environment data that is obtained from the ISE when you complete the following tasks on the ASA:

- Configure a AAA server to communicate with the ISE.
- Import a PAC file from the ISE.
- Identify the AAA server group that the ASA will use to retrieve Cisco TrustSec environment data.

Normally, you do not need to manually refresh the environment data from the ISE; however, security groups can change on the ISE. These changes are not reflected on the ASA until you refresh the data in the ASA

security group table, so refresh the data on the ASA to make sure that any security group changes made on the ISE are reflected on the ASA.

Note We recommend that you schedule policy configuration changes on the ISE and the manual data refresh on the ASA during a maintenance window. Handling policy configuration changes in this way maximizes the chances of security group names getting resolved and security policies becoming active immediately on the ASA.

To refresh the environment data, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Identity By TrustSec**.

Step 2 Click **Refresh Environment** > **Data** in the **Server Group Setup** area.

The ASA refreshes the Cisco TrustSec environment data from the ISE and resets the reconcile timer to the configured default value.

Configure the Security Policy

You can incorporate Cisco TrustSec policy in many ASA features. Any feature that uses extended ACLs (unless listed in this chapter as unsupported) can take advantage of Cisco TrustSec. You can add security group arguments to extended ACLs, as well as traditional network-based parameters.

- To configure access rules, see Configure Access Rules, on page 19. For other extended ACLs, see Configure Extended ACLs, on page 47.
- To configure security group object groups that can be used in the ACL, see Configure Security Group Object Groups, on page 36.

For example, an access rule permits or denies traffic on an interface using network information. With Cisco TrustSec, you can control access based on security group. For example, you could create an access rule for sample_securitygroup1 10.0.0.0 255.0.0.0, meaning the security group could have any IP address on subnet 10.0.0.0/8.

You can configure security policies based on combinations of security group names (servers, users, unmanaged devices, and so on), user-based attributes, and traditional IP-address-based objects (IP address, Active Directory object, and FQDN). Security group membership can extend beyond roles to include device and location attributes and is independent of user group membership.

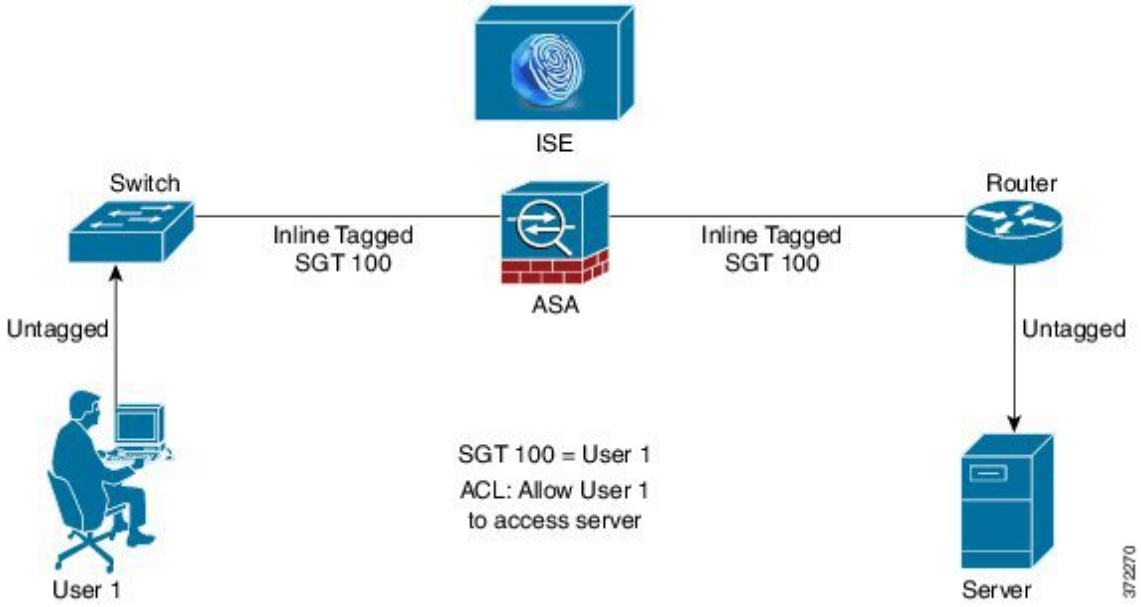
Configure Layer 2 Security Group Tagging Imposition

Cisco TrustSec identifies and authenticates each network user and resource and assigns a 16-bit number called a Security Group Tag (SGT). This identifier is in turn propagated between network hops, which allows any intermediary devices such as ASAs, switches, and routers to enforce polices based on this identity tag.

SGT plus Ethernet Tagging, also called Layer 2 SGT Imposition, enables the ASA to send and receive security group tags on Ethernet interfaces using Cisco proprietary Ethernet framing (EtherType 0x8909), which allows the insertion of source security group tags into plain-text Ethernet frames. The ASA inserts security group tags on the outgoing packet and processes security group tags on the incoming packet, based on a manual per-interface configuration. This feature allows inline hop-by-hop propagation of endpoint identity across network devices and provides seamless Layer 2 SGT Imposition between each hop.

The following figure shows a typical example of Layer 2 SGT Imposition.

Figure 12: Layer 2 SGT Imposition



Usage Scenarios

The following table describes the expected behavior for ingress traffic when configuring this feature.

Table 4: Ingress Traffic

Interface Configuration	Tagged Packet Received	Untagged Packet Received
No command is issued.	Packet is dropped.	SGT value is from the IP-SGT Manager.
The **cts manual** command is issued.	SGT value is from the IP-SGT Manager.	SGT value is from the IP-SGT Manager.
The **cts manual** command and the **policy static sgt** *sgt_number* command are both issued.	SGT value is from the **policy static sgt** *sgt_number* command.	SGT value is from the **policy static sgt** *sgt_number* command.
The **cts manual** command and the **policy static sgt** *sgt_number* **trusted** command are both issued.	SGT value is from the inline SGT in the packet.	SGT value is from the **policy static sgt** *sgt_number* command.

Note If there is no matched IP-SGT mapping from the IP-SGT Manager, then a reserved SGT value of "0x0" for "Unknown" is used.

The following table describes the expected behavior for egress traffic when configuring this feature.

Table 5: Egress Traffic

Interface Configuration	Tagged or Untagged Packet S
No command is issued.	Untagged
The **cts manual** command is issued.	Tagged
The **cts manual** command and the **propagate sgt** command are both issued.	Tagged
The **cts manual** command and the **no propagate sgt** command are both issued.	Untagged

The following table describes the expected behavior for to-the-box and from-the-box traffic when configuring this feature.

Table 6: To-the-box and From-the-box Traffic

Interface Configuration	Tagged or Untagged Packet Received
No command is issued on the ingress interface for to-the-box traffic.	Packet is dropped.
The **cts manual** command is issued on the ingress interface for to-the-box traffic.	Packet is accepted, but there is no policy enforcem propagation.
The **cts manual** command is not issued or the **cts manual** command and **no propagate sgt** command are both issued on the egress interface for from-the-box traffic.	Untagged packet is sent, but there is no policy enfo SGT number is from the IP-SGT Manager.
The **cts manual** command is issued or the **cts manual** command and the **propagate sgt** command are both issued on the egress interface for from-the-box traffic.	Tagged packet is sent. The SGT number is from th Manager.

Note If there is no matched IP-SGT mapping from the IP-SGT Manager, then a reserved SGT value of "0x0" for "Unknown" is used.

Configure a Security Group Tag on an Interface

To configure a security group tag on an interface, perform the following steps:

Procedure

Step 1 Choose one of the following options:

- **Configuration > Device Setup** > **Interfaces> Add Interface > Advanced**
- **Configuration > Device Setup** > **Interfaces > Add Redundant Interface > Advanced**
- **Configuration > Device Setup** > **Interfaces > Add Ethernet Interface > Advanced**

Step 2 Check the **Enable secure group tagging for Cisco TrustSec** check box.

Step 3 Check the **Tag egress packets with service group tags** check box.

Step 4 Check the **Add a static secure group tag to all ingress packets** check box.

Step 5 Enter a secure group tag number. Valid values range from 2 - 65519.

Step 6 Check the **This is a trusted interface. Do not override existing secure group tags** check box.

Step 7 Click **OK** to save your settings.

Configure IP-SGT Bindings Manually

To configure IP-SGT bindings manually, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Firewall Identity by TrustSec**.

Step 2 Click **Add** in the **SGT Map Setup** area, or select an SGT map and click **Edit**.

Step 3 In the SGT Map dialog box, enter the SGT Map IP address and the SGT value in the appropriate fields.

SGT numbers can be from 2 to 65519.

To map a network to an SGT, select the **Prefix** check box and enter the subnet or IPv6 prefix. For example, enter 24 to map 10.100.10.0/24.

Step 4 Click **OK**, then click **Apply** to save your settings.

AnyConnect Client VPN Support for Cisco TrustSec

ASA supports security group tagging of VPN sessions. You can assign a Security Group Tag (SGT) to a VPN session using an external AAA server, or by configuring a security group tag for a local user or for a VPN group policy. This tag can then be propagated through the Cisco TrustSec system over Layer 2 Ethernet. Security group tags are useful on group policies and for local users when the AAA server cannot provide an SGT.

Following is the typical process for assigning an SGT to a VPN user:

1. A user connects to a remote access VPN that uses a AAA server group containing ISE servers.

2. The ASA requests AAA information from ISE, which might include an SGT. The ASA also assigns an IP address for the user's tunneled traffic.
3. The ASA uses AAA information to authenticate the user and creates a tunnel.
4. The ASA uses the SGT from AAA information and the assigned IP address to add an SGT in the Layer 2 header.
5. Packets that include the SGT are passed to the next peer device in the Cisco TrustSec network.

If there is no SGT in the attributes from the AAA server to assign to a VPN user, then the ASA uses the SGT in the group policy. If there is no SGT in the group policy, then tag 0x0 is assigned.



Note You can also use ISE for policy enforcement using ISE Change of Authorization (CoA). For information on how to configure policy enforcement, see the VPN configuration guide.

Add an SGT to Remote Access VPN Group Policies and Local Users

To configure an SGT attribute on remote access VPN group policies, or on the VPN policy for a user defined in the LOCAL user database, perform the following steps.

There is no default SGT for group policies or local users.

Procedure

Step 1 To configure an SGT on a remote access VPN group policy:

a) Choose **Configuration** > **Remote Access VPN** > **Network (Client) Access** > **Group Policies**.
b) Click the **General** tab, then click **More Options**.
c) Enter a value in the **Security Group Tag (STG)** field, from 2 to 65519.

You can also select None to set no SGT.

d) Click **OK**.

Step 2 To configure an SGT on for a user in the LOCAL database:

a) Choose **Configuration** > **Remote Access VPN** > **AAA/Local Users** > **Local Users**.
b) Select a user, then click **Edit**.
c) Click **VPN Policy**.
d) Enter a value in the **Security Group Tag (STG)** field, from 2 to 65519.

You can also select None to set no SGT.

e) Click **OK**.

Monitoring Cisco TrustSec

See the following screens for monitoring Cisco TrustSec:

- **Monitoring** > **Properties** > **Identity By TrustSec** > **SXP Connections**

  Shows the configured default values for the Cisco TrustSec infrastructure and the SXP commands.

- **Monitoring** > **Properties** > **Connections**

  Filters the IP address-security group table mapping entries so that you view the data by security group table value, security group name, or IP address.

- **Monitoring** > **Properties** > **Identity By TrustSec** > **Environment Data**

  Shows the Cisco TrustSec environment information contained in the security group table on the ASA.

- **Monitoring** > **Properties** > **Identity By TrustSec** > **IP Mapping**

  Filters the IP address-security group table mapping entries so that you view the data by security group table value, security group name, or IP address. Click **Where Used** to show where the selected security group object is used in an ACL or nested in another security group object.

- **Monitoring** > **Properties** > **Identity By TrustSec** > **PAC**

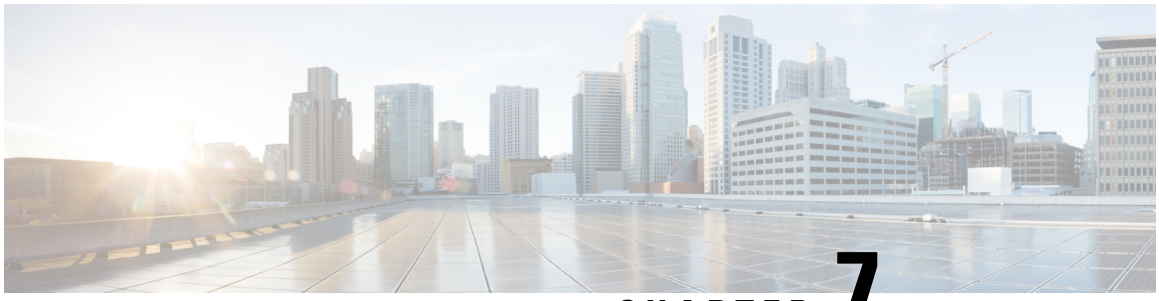
  Shows information about the PAC file imported into the ASA from the ISE and includes a warning message when the PAC file has expired or is within 30 days of expiration.

History for Cisco TrustSec

Table 7: History for Cisco TrustSec

Feature Name	Platform Releases	Description
Cisco TrustSec	9.0(1)	Cisco TrustSec provides access control that builds on an existing identity-aware infrastructure to ensure data confidentiality between network devices and integrate security access services on one platform. In the Cisco TrustSec feature, enforcement devices use a combination of user attributes and endpoint attributes to make role-based and identity-based access control decisions. In this release, the ASA integrates with Cisco TrustSec to provide security group-based policy enforcement. Access policies within the Cisco TrustSec domain are topology-independent, based on the roles of source and destination devices rather than on network IP addresses. The ASA can use Cisco TrustSec for other types of security group-based policies, such as application inspection; for example, you can configure a class map that includes an access policy based on a security group. We introduced or modified the following screens: Configuration > Firewall > Identity By TrustSec Configuration > Firewall > Objects > Security Groups Object Groups Configuration > Firewall > Access Rules > Add Access Rules Monitoring > Properties > Identity By Tag.
Layer 2 Security Group Tag Imposition	9.3(1)	You can now use security group tagging combined with Ethernet tagging to enforce policies. SGT plus Ethernet Tagging, also called Layer 2 SGT Imposition, enables the ASA to send and receive security group tags on Ethernet interfaces using Cisco proprietary Ethernet framing (EtherType 0x8909), which allows the insertion of source security group tags into plain-text Ethernet frames. We modified the following screens: Configuration > Device Setup > Interfaces > Add Interface > Advanced Configuration > Device Setup > Interfaces > Add Redundant Interface > Advanced Configuration > Device Setup > Add Ethernet Interface > Advanced.
Cisco Trustsec support for Security Exchange Protocol (SXP) version 3.	9.6(1)	Cisco Trustsec on ASA now implements SXPv3, which enables SGT-to-subnet bindings, which are more efficient than host bindings. We modified the following screens: **Configuration** > **Firewall** > **Identity By TrustSec** and the **SGT Map Setup** dialog boxes.

Feature Name	Platform Releases	Description
Trustsec SXP connection configurable delete hold down timer	9.8(3)	The default SXP connection hold down timer is 120 seconds. You can now configure this timer, between 120 to 64000 seconds. New/Modified commands: **cts sxp delete-hold-down period**, **show cts sxp connection brief**, **show cts sxp connections** No ASDM support.



CHAPTER 7

ASA FirePOWER Module

The following topics describe how to configure the ASA FirePOWER module that runs on the ASA.

- About the ASA FirePOWER Module, on page 97
- Licensing Requirements for the ASA FirePOWER Module, on page 101
- Guidelines for ASA FirePOWER, on page 101
- Defaults for ASA FirePOWER, on page 103
- Perform Initial ASA FirePOWER Setup, on page 103
- Configure the ASA FirePOWER Module, on page 111
- Managing the ASA FirePOWER Module, on page 114
- Monitoring the ASA FirePOWER Module, on page 121
- History for the ASA FirePOWER Module, on page 123

About the ASA FirePOWER Module

The ASA FirePOWER module supplies next-generation firewall services, including Next-Generation Intrusion Prevention System (NGIPS), Application Visibility and Control (AVC), URL filtering, and Advanced Malware Protection (AMP).

The ASA FirePOWER module runs a separate application from the ASA.

How the ASA FirePOWER Module Works with the ASA

You can configure your ASA FirePOWER module using one of the following deployment models:

- Inline mode—In an inline deployment, the actual traffic is sent to the ASA FirePOWER module, and the module's policy affects what happens to the traffic. After dropping undesired traffic and taking any other actions applied by policy, the traffic is returned to the ASA for further processing and ultimate transmission.
- Inline tap monitor-only mode (ASA inline)—In an inline tap monitor-only deployment, a copy of the traffic is sent to the ASA FirePOWER module, but it is not returned to the ASA. Inline tap mode lets you see what the ASA FirePOWER module would have done to traffic, and lets you evaluate the content of the traffic, without impacting the network. However, in this mode, the ASA does apply its policies to the traffic, so traffic can be dropped due to access rules, TCP normalization, and so forth.
- Passive monitor-only (traffic forwarding) mode—If you want to prevent any possibility of the ASA with FirePOWER Services device impacting traffic, you can configure a traffic-forwarding interface and

connect it to a SPAN port on a switch. In this mode, traffic is sent directly to the ASA FirePOWER module without ASA processing. The traffic is dropped, and nothing is returned from the module, nor does the ASA send the traffic out any interface. You must operate the ASA in single context transparent mode to configure traffic forwarding.

Be sure to configure consistent policies on the ASA and the ASA FirePOWER. Both policies should reflect the inline or monitor-only mode of the traffic.

The following sections explain these modes in more detail.

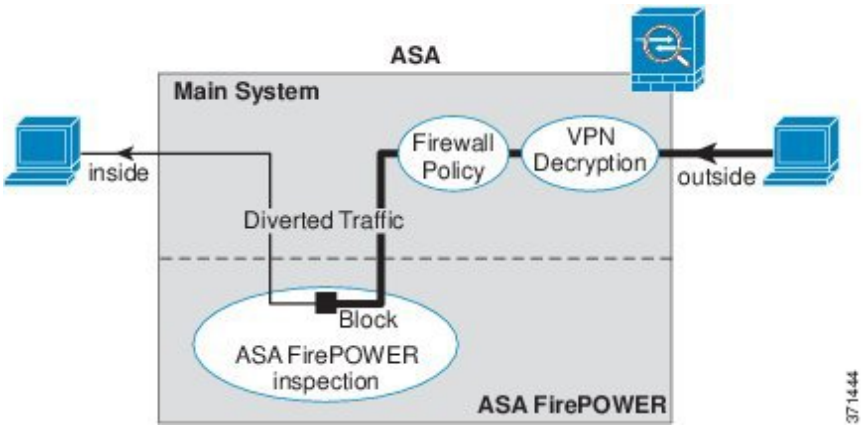
ASA FirePOWER Inline Mode

In inline mode, traffic goes through the firewall checks before being forwarded to the ASA FirePOWER module. When you identify traffic for ASA FirePOWER inspection on the ASA, traffic flows through the ASA and the module as follows:

1. Traffic enters the ASA.
2. Incoming VPN traffic is decrypted.
3. Firewall policies are applied.
4. Traffic is sent to the ASA FirePOWER module.
5. The ASA FirePOWER module applies its security policy to the traffic, and takes appropriate actions.
6. Valid traffic is sent back to the ASA; the ASA FirePOWER module might block some traffic according to its security policy, and that traffic is not passed on.
7. Outgoing VPN traffic is encrypted.
8. Traffic exits the ASA.

The following figure shows the traffic flow when using the ASA FirePOWER module in inline mode. In this example, the module blocks traffic that is not allowed for a certain application. All other traffic is forwarded through the ASA.

Figure 13: ASA FirePOWER Module Traffic Flow in the ASA



Note If you have a connection between hosts on two ASA interfaces, and the ASA FirePOWER service policy is only configured for one of the interfaces, then all traffic between these hosts is sent to the ASA FirePOWER module, including traffic originating on the non-ASA FirePOWER interface (because the feature is bidirectional).

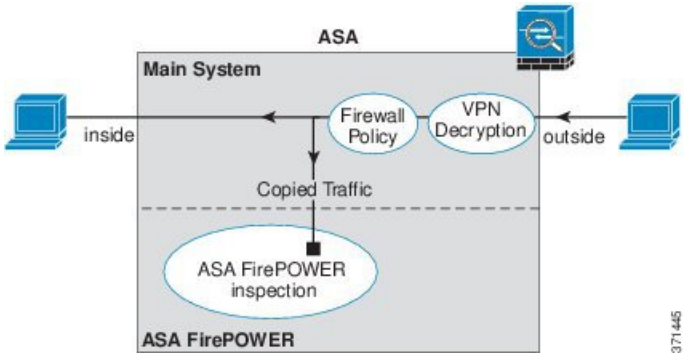
ASA FirePOWER Inline Tap Monitor-Only Mode

This mode sends a duplicate stream of traffic to the ASA FirePOWER module for monitoring purposes only. The module applies the security policy to the traffic and lets you know what it would have done if it were operating in inline mode; for example, traffic might be marked "would have dropped" in events. You can use this information for traffic analysis and to help you decide if inline mode is desirable.

Note You cannot configure both inline tap monitor-only mode and normal inline mode at the same time on the ASA. Only one type of service policy rule is allowed. In multiple context mode, you cannot configure inline tap monitor-only mode for some contexts, and regular inline mode for others.

The following figure shows the traffic flow when operating in inline tap mode.

Figure 14: ASA FirePOWER Inline Tap Monitor-Only Mode



ASA FirePOWER Passive Monitor-Only Traffic Forwarding Mode

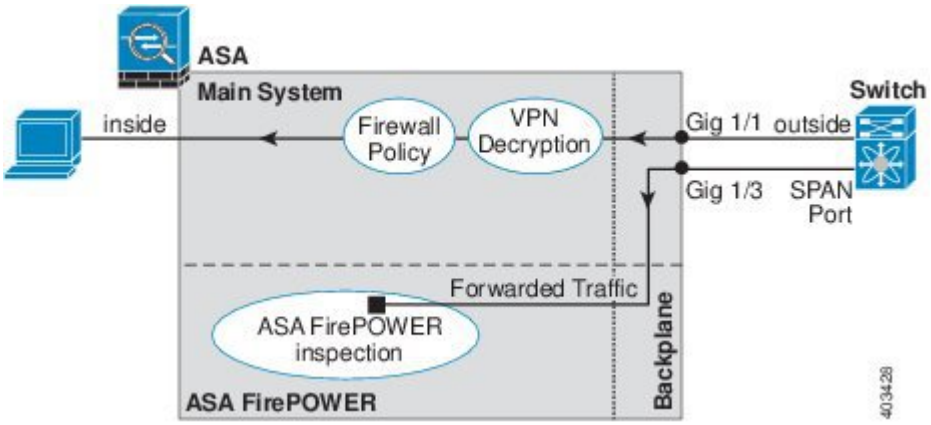
If you want to operate the ASA FirePOWER module as a pure Intrusion Detection System (IDS), where there is no impact on the traffic at all, you can configure a traffic forwarding interface. A traffic forwarding interface sends all received traffic directly to the ASA FirePOWER module without any ASA processing.

The module applies the security policy to the traffic and lets you know what it would have done if it were operating in inline mode; for example, traffic might be marked "would have dropped" in events. You can use this information for traffic analysis and to help you decide if inline mode is desirable.

Traffic in this setup is never forwarded: neither the module nor the ASA sends the traffic on to its ultimate destination. You must operate the ASA in single context and transparent modes to use this configuration.

The following figure shows an interface configured for traffic-forwarding. That interface is connected to a switch SPAN port so the ASA FirePOWER module can inspect all of the network traffic. Another interface sends traffic normally through the firewall.

Figure 15: ASA FirePOWER Passive Monitor-Only, Traffic-Forwarding Mode



ASA FirePOWER Management

The module has a basic command line interface (CLI) for initial configuration and troubleshooting only. You configure the security policy on the ASA FirePOWER module using one of the following methods:

- Firepower/FireSIGHT Management Center—Can be hosted on a separate Management Center appliance or as a virtual appliance. The Management Center application is called Firepower beginning in version 6.0. Previous versions are called FireSIGHT.
- ASDM (check for compatibility with your model/version)—You can manage both the ASA and the module using the on-box ASDM.

Compatibility with ASA Features

The ASA includes many advanced application inspection features, including HTTP inspection. However, the ASA FirePOWER module provides more advanced HTTP inspection than the ASA provides, as well as additional features for other applications, including monitoring and controlling application usage.

You must follow these configuration restrictions on the ASA:

- Do not configure ASA inspection on HTTP traffic that you send to the ASA FirePOWER module.
- Do not enable the Mobile User Security (MUS) server; it is not compatible with the ASA FirePOWER module.

Other application inspections on the ASA are compatible with the ASA FirePOWER module, including the default inspections.

What to Do if the ASA FirePOWER Module Cannot Filter URLs

The ASA FirePOWER module obtains its URL filtering data from the managing Firepower Management Center, over HTTP. The module cannot perform URL filtering if it cannot download this database.

If there is a device between the ASA FirePOWER module and Firepower Management Center that is performing ASA HTTP inspection, the inspections can block the ASA FirePOWER module's HTTP GET requests to the

Firepower Management Center. This problem also occurs if you configure HTTP inspection on the ASA that hosts the ASA FirePOWER module (which is a misconfiguration).

To resolve the issues, do any of the following that apply to your situation:

- If you configured HTTP inspection on the ASA that hosts the ASA FirePOWER module, remove the HTTP inspection configuration. ASA FirePOWER inspection and ASA HTTP inspection are incompatible.
- If there is an intervening device doing ASA HTTP inspection, remove the drop protocol violations action from the HTTP inspection policy map:

```
policy-map type inspect http http_inspection_policy
   parameters
      no protocol-violation action drop-connection
```

Licensing Requirements for the ASA FirePOWER Module

Certain areas of ASA FirePOWER module functionality may require additional licenses.

For an ASA FirePOWER module managed by a Firepower/FireSIGHT Management Center, enable licenses on the module using the Management Center. See the licensing chapter of the *FireSIGHT System User Guide 5.4*, *Firepower Management Center Configuration Guide 6.0*, or the online help on the FireSIGHT Management Center for more information.

For the ASA FirePOWER module managed using ASDM, enable licenses on the module using the FirePOWER module configuration in ASDM. See the licensing chapter of the *ASA FirePOWER Module User Guide 5.4*, *ASA FirePOWER Services Local Management Configuration Guide 6.0*, or the online help for the module in ASDM for more information.

The ASA itself does not require any additional licenses.

Guidelines for ASA FirePOWER

Failover Guidelines

Does not support failover directly; when the ASA fails over, any existing ASA FirePOWER flows are transferred to the new ASA. The ASA FirePOWER module in the new ASA begins inspecting the traffic from that point forward; old inspection states are not transferred.

You are responsible for maintaining consistent policies on the ASA FirePOWER modules in the high-availability ASA pair to ensure consistent failover behavior.



Note Create the failover pair before you configure the ASA FirePOWER modules. If the modules are already configured on both devices, clear the interface configuration on the standby device before creating the high-availability pair. From the CLI on the standby device, enter the **clear configure interface** command.

ASA Clustering Guidelines

Does not support clustering directly, but you can use these modules in a cluster. You are responsible for maintaining consistent policies on the ASA FirePOWER modules in the cluster.



Note Create the cluster before you configure the ASA FirePOWER modules. If the modules are already configured on the data units, clear the interface configuration on the devices before adding them to the cluster. From the CLI, enter the **clear configure interface** command.

Model Guidelines

- For ASA model software and hardware compatibility with the ASA FirePOWER module, see the Cisco ASA Compatibility.
- For the ASA 5515-X through ASA 5555-X, you must install a Cisco solid state drive (SSD). For more information, see the ASA 5500-X hardware guide. (The SSD is standard on the 5508-X, and 5516-X.)

ASDM Guidelines for Managing ASA FirePOWER

- The ASA, ASDM, and ASA FirePOWER versions supported for ASDM management differ by model. For supported combinations, see Cisco ASA Compatibility.
- If you enable command authorization on the ASA that hosts the module, you must log in with a user name that has privilege level 15 to see the **ASA FirePOWER** home, configuration, and monitoring pages. Read-only or monitor-only access to **ASA FirePOWER** pages other than the status page is not supported.
- If you are using Java 7 update 51 up to Java 8, you need to configure identity certificates for both the ASA and the ASA FirePOWER module. See Install an Identity Certificate for ASDM.
- You can never use both ASDM and Firepower Management Center, you must choose one or the other.

Additional Guidelines and Limitations

- See Compatibility with ASA Features, on page 100.
- You cannot configure both normal inline mode and inline tap monitor-only mode at the same time on the ASA. Only one type of service policy rule is allowed. In multiple context mode, you cannot configure inline tap monitor-only mode for some contexts, and regular inline mode for others.
- If you configure NetFlow on the ASA and include the **flow-export delay flow-create** command, even if your ASA FirePOWER access control policy blocks the connection with reset, the connection will remain on the ASA until the connection timeout is reached. If you cannot tolerate this behavior, you need to remove the command from the NetFlow configuration.
- If the module is stuck in recover/init mode, you cannot reload the ASA gracefully. Instead, use the **reload quick** command, so the ASA does not wait for the module to shut down gracefully prior to the system reload. If quick reload does not work, you will have to force-crash the ASA to reload it.
- On the ASA 5500-X series, particularly the smaller models, you can see very intermittent latency, 30-60ms, on .02% of the traffic. If you have applications that cannot tolerate this small latency, ensure

that you configure your redirection service policy so that these latency-sensitive applications are not redirected to the ASA FirePOWER module.

Defaults for ASA FirePOWER

The following table lists the default settings for the ASA FirePOWER module.

Table 8: ASA FirePOWER Default Network Parameters

Parameters	Default
Management IP address	System software image: 192.168.45.45/24 Boot image: 192.168.8.8/24
Gateway	System software image: none Boot image: 192.168.8.1/24
SSH or session Username	admin
Password	System software image: • ISA 3000 only on a new factory default image in Release 7.0.1: **Adm!n123** • Release 6.0 and following: **Admin123** • Releases prior to 6.0: **Sourcefire** Boot image: **Admin123**

Perform Initial ASA FirePOWER Setup

Deploy the ASA FirePOWER module in your network, and then choose your management method.

Deploy the ASA FirePOWER Module in Your Network

See the section for your firewall mode and ASA model to determine how to connect the ASA FirePOWER module management interface to your network.

Routed Mode

ASA 5508-X through ASA 5555-X (Software Module) in Routed Mode

These models run the ASA FirePOWER module as a software module, and the ASA FirePOWER module shares the Management 0/0 or Management 1/1 interface (depending on your model) with the ASA.

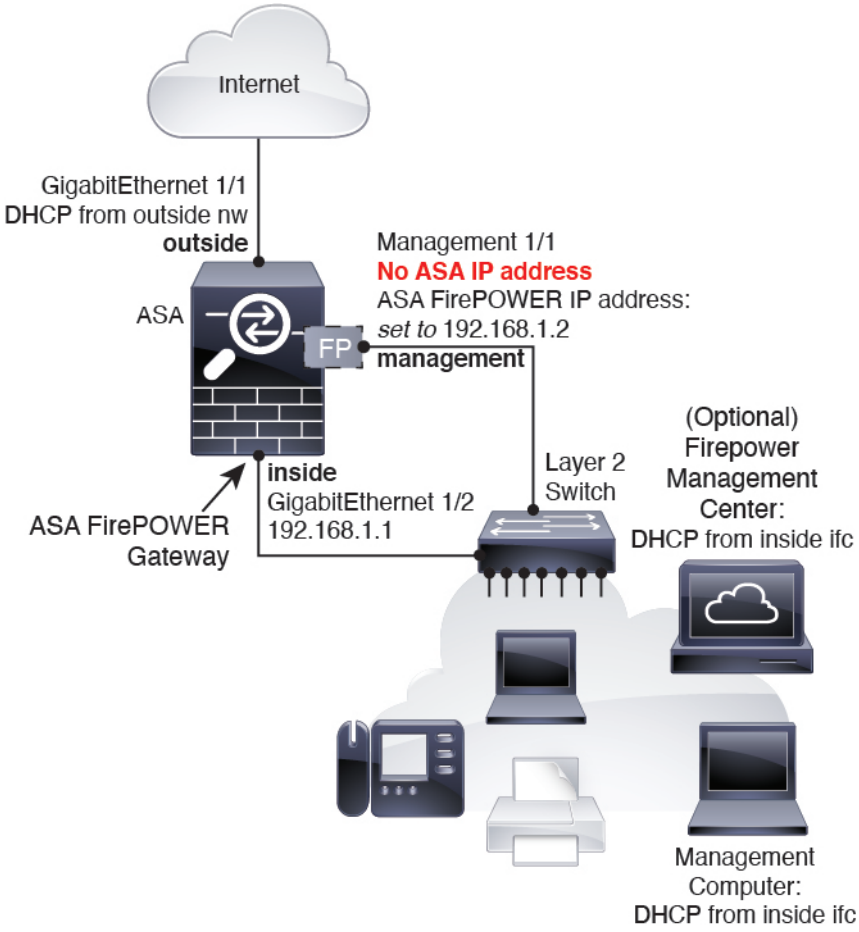
All management traffic to and from the ASA FirePOWER module must enter and exit the Management interface. The ASA FirePOWER module also needs Internet access. Management traffic cannot pass through

the ASA over the backplane; therefore you need to physically cable the management interface to an ASA interface to reach the Internet.

If you do not configure a name and IP address in the ASA configuration for Management, then the interface belongs exclusively to the module. In this case, the Management interface is not a regular ASA interface, and you can:

1. Configure the ASA FirePOWER IP address to be on the same network as a regular ASA data interface.
2. Specify the data interface as the ASA FirePOWER gateway.
3. Directly connect the Management interface to the data interface (using a Layer2 switch).

See the following typical cabling setup to allow ASA FirePOWER access to the Internet through the ASA inside interface.



For the ASA 5508-X, and 5516-X, the default configuration enables the above network deployment; the only change you need to make is to set the module IP address to be on the same network as the ASA inside interface and to configure the module gateway IP address.

For other models, you must remove the ASA-configured name and IP address for Management 0/0 or 1/1, and then configure the other interfaces as indicated above.

Note You can avoid using an external switch if you have extra interfaces that you can assign to an inside bridge group to configure a "soft switch". Be sure to set all bridge group interfaces to the same security level, allow same security communication, and configure NAT for each bridge group member. See the ASA interfaces configuration guide chapter for more information.

Note If you want to deploy a separate router on the inside network, then you can route between management and inside. In this case, you can manage both the ASA and ASA FirePOWER module on the Management interface with the appropriate configuration changes, including configuring the ASA name and IP address for the Management interface (on the same network as the ASA FirePOWER module address).

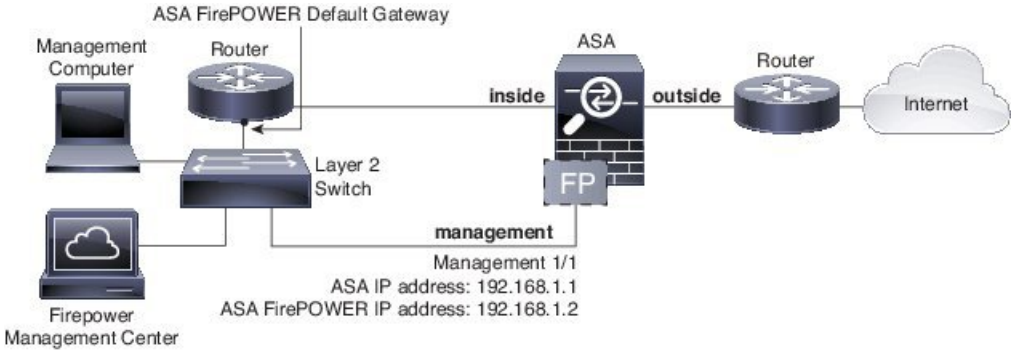
Transparent Mode

ASA 5508-X through ASA 5555-X, ISA 3000 (Software Module) in Transparent Mode

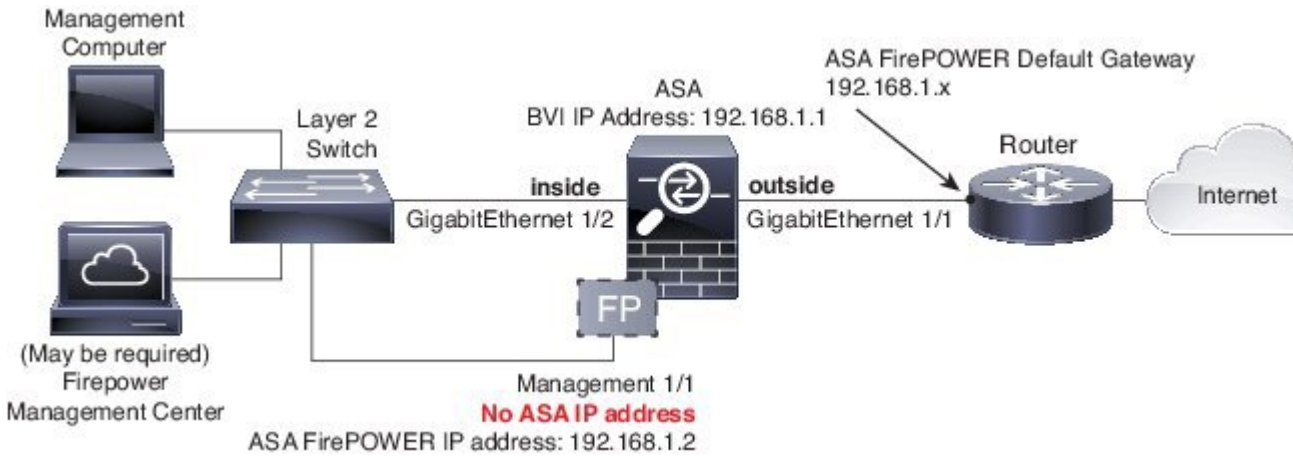
These models run the ASA FirePOWER module as a software module, and the ASA FirePOWER module shares the Management 0/0 or Management 1/1 interface (depending on your model) with the ASA.

All management traffic to and from the ASA FirePOWER module must enter and exit the Management interface. The ASA FirePOWER module also needs Internet access.

The following figure shows the recommended network deployment for the ASA 5500-X or ISA 3000 with the ASA FirePOWER module:



If you do not use an inside router, you can manage the ASA over the inside interface (using the BVI IP address) and not use the Management interface for ASA management:



Note You can avoid using an external switch if you have extra interfaces that you can assign to the inside bridge group to configure a "soft switch". Be sure to set all bridge group interfaces to the same security level, allow same security communication, and configure NAT for each bridge group member. See the ASA interfaces configuration guide chapter for more information.

Register the ASA FirePOWER Module with a Management Center

To register the module with a Firepower/FireSIGHT Management Center, you must access the ASA FirePOWER module CLI. The first time you access the CLI, you are prompted for basic configuration parameters. You must also add the module to the Management Center.

Notes:

- If you want to use ASDM to manage the module, skip this section and see Configure the ASA FirePOWER Module for ASDM Management, on page 108.
- If you need to move the module's management from one management center to another, first remove the device from the management center's inventory. Then, use the **configure manager add** command to point to the new management center. You can then complete the registration from the new management center. This process ensures a clean hand-over.

Access the ASA FirePOWER CLI

To access the ASA FirePOWER CLI, you can use one of the following methods.

Procedure

Step 1 Console Port:

- Connect to the ASA console port using the supplied DB-9 to RJ-45 serial cable and/or your own USB serial adapter. The ASA 5508-X/5516-X also has a mini-USB console port. See the hardware guide for instructions on using the USB console port.

At the ASA CLI, session to the ASA FirePOWER module:

session sfr

See also Session to the Software Module From the ASA, on page 120.

Step 2 SSH:

You can connect to the module default IP address (see Defaults for ASA FirePOWER, on page 103) or you can use ASDM on the ASA to change the management IP address, and then connect using SSH:

In ASDM, choose **Wizards > Startup Wizard**, and progress through the wizard to the **ASA FirePOWER Basic Configuration**, where you can set the IP address, mask, and default gateway.

Configure ASA FirePOWER Basic Settings

The first time you access the ASA FirePOWER module CLI, you are prompted for basic configuration parameters. You must also add the module to the Firepower/FireSIGHT Management Center if you are not using ASDM.

Before you begin

Access the module CLI according to Access the ASA FirePOWER CLI, on page 106.

Procedure

Step 1 At the ASA FirePOWER CLI, log in with the username **admin**.

If this is the first time you are logging in, use the default password. See Defaults for ASA FirePOWER, on page 103.

Step 2 Complete the system configuration as prompted.

Use the following network settings for the ASA FirePOWER module for the recommended network deployment (Deploy the ASA FirePOWER Module in Your Network, on page 103):

- Management interface: 192.168.1.2
- Management subnet mask: 255.255.255.0
- Gateway IP: 192.168.1.1

Example:

```
System initialization in progress.  Please stand by.
You must change the password for 'admin' to continue.
Enter new password: <new password>
Confirm new password: <repeat password>
You must configure the network to continue.
You must configure at least one of IPv4 or IPv6.
Do you want to configure IPv4? (y/n) [y]: y
Do you want to configure IPv6? (y/n) [n]:
Configure IPv4 via DHCP or manually? (dhcp/manual) [manual]:
Enter an IPv4 address for the management interface [192.168.45.45]: 10.86.118.3
```

```
Enter an IPv4 netmask for the management interface [255.255.255.0]: 255.255.252.0
Enter the IPv4 default gateway for the management interface []: 10.86.116.1
Enter a fully qualified hostname for this system [Sourcefire3D]: asasfr.example.com
Enter a comma-separated list of DNS servers or 'none' []: 10.100.10.15,
10.120.10.14
Enter a comma-separated list of search domains or 'none' [example.net]: example.com
If your networking information has changed, you will need to reconnect.
For HTTP Proxy configuration, run 'configure network http-proxy'
(Wait for the system to reconfigure itself.)

This sensor must be managed by a Defense Center.  A unique alphanumeric
registration key is always required.  In most cases, to register a sensor
to a Defense Center, you must provide the hostname or the IP address along
with the registration key.
'configure manager add [hostname | ip address ] [registration key ]'

However, if the sensor and the Defense Center are separated by a NAT device,
you must enter a unique NAT ID, along with the unique registration key.
'configure manager add DONTRESOLVE [registration key ] [ NAT ID ]'

Later, using the web interface on the Defense Center, you must use the same
registration key and, if necessary, the same NAT ID when you add this
sensor to the Defense Center.
```

Step 3 Register the ASA FirePOWER module to a Management Center:

> **configure manager add** {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE**} *reg_key* [*nat_id*]

where:

- {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE**} specifies either the fully qualified host name or IP address of the Management Center. If the Management Center is not directly addressable, use DONTRESOLVE.
- *reg_key* is the unique alphanumeric registration key required to register a ASA FirePOWER module to the Management Center.
- *nat_id* is an optional alphanumeric string used during the registration process between the Management Center and the ASA FirePOWER module. It is required if the hostname is set to DONTRESOLVE.

Step 4 Close the console connection. For the software module, enter:

> **exit**

Configure the ASA FirePOWER Module for ASDM Management

Not supported for all version/model combinations; check for compatibility with your model and version.

ASDM can change the ASA FirePOWER module IP address over the ASA backplane, but all further management requires network access between the ASDM interface and the Management interface, where the module is reachable.

To use ASDM to manage the module, launch ASDM and run the Startup Wizard.

Procedure

Step 1 On the computer connected to the ASA, launch a web browser.

Step 2 In the Address field, enter the following URL: **https://192.168.1.1/admin**. The Cisco ASDM web page appears.

Step 3 Click one of the available options: **Install ASDM Launcher**, **Run ASDM**, or **Run Startup Wizard**.

Step 4 Follow the onscreen instructions to launch ASDM according to the option you chose. The Cisco ASDM-IDM Launcher appears.

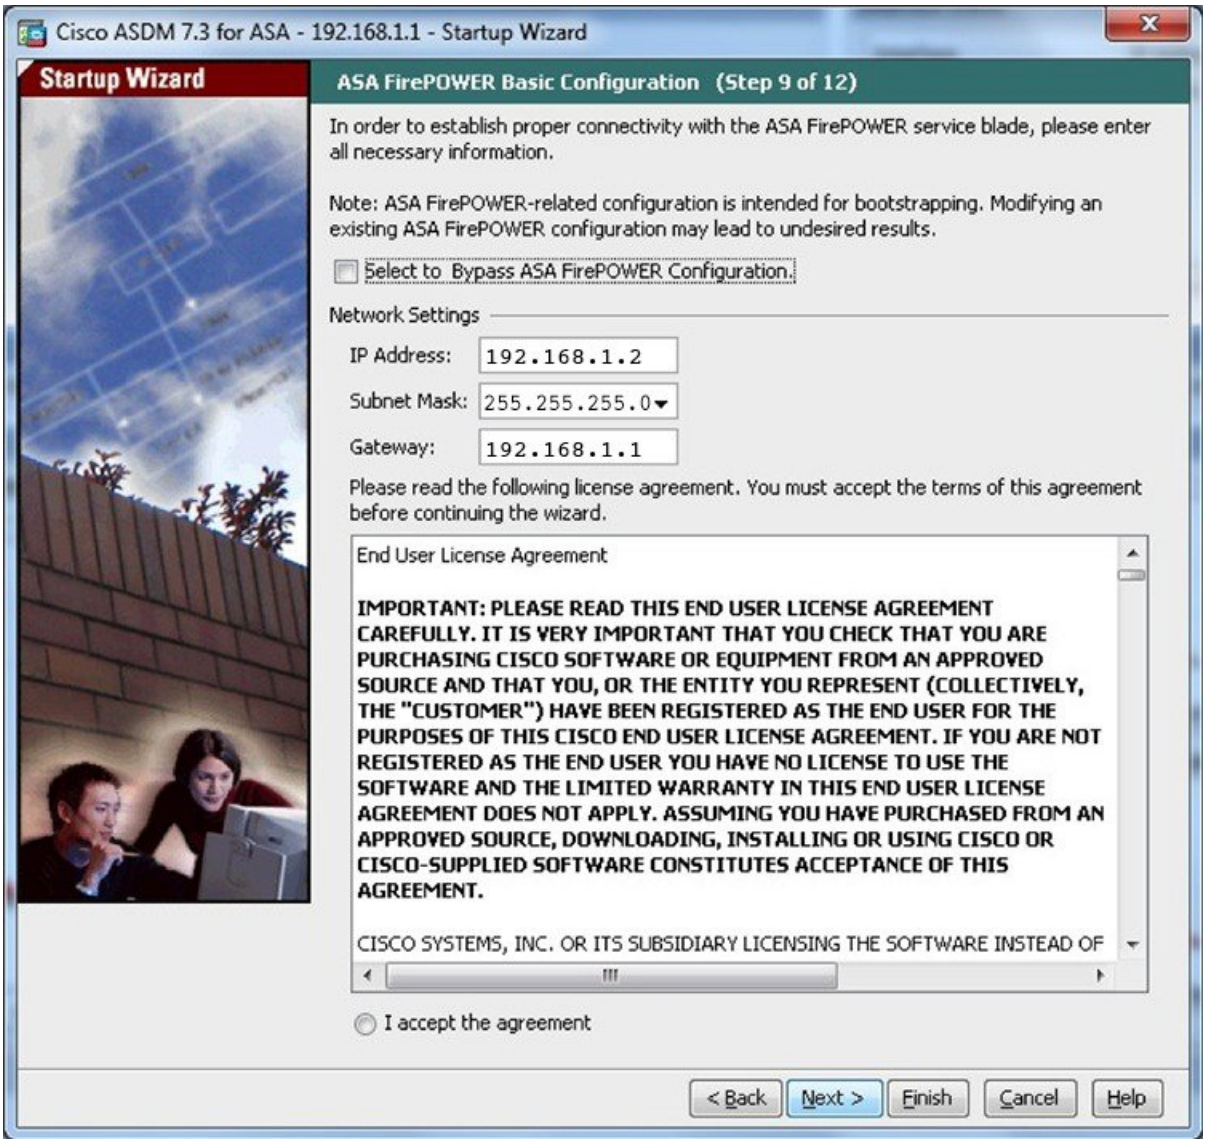
Note
If you click Install ASDM Launcher, in some cases you need to install an identity certificate for the ASA and a separate certificate for the ASA FirePOWER module according to Install an Identity Certificate for ASDM.

Step 5 Leave the username and password fields empty, and click **OK**. The main ASDM window appears.

Step 6 If you are prompted to provide the IP address of the installed ASA Firepower module, cancel out of the dialog box. You must first set the module IP address to the correct IP address using the Startup Wizard.

Step 7 Choose **Wizards** > **Startup Wizard**.

Step 8 Configure additional ASA settings as desired, or skip screens until you reach the **ASA Firepower Basic Configuration** screen.



Set the following values to work with the default configuration:

- **IP Address**—192.168.1.2
- **Subnet Mask**—255.255.255.0
- **Gateway**—192.168.1.1

Step 9 Click I accept the agreement, and click **Next** or **Finish** to complete the wizard.

Step 10 Quit ASDM, and then relaunch. You should see **ASA Firepower** tabs on the **Home** page.

Configure the ASA FirePOWER Module

Configure the security policy in the ASA FirePOWER module, and then configure the ASA to send traffic to the module.

Configure the Security Policy on the ASA FirePOWER Module

The security policy controls the services provided by the module, such as Next Generation IPS filtering and application filtering. You configure the security policy on the ASA FirePOWER module using one of the following methods.

FireSIGHT Management Center

Use a web browser to open **https://***DC_address*, where *DC_address* is the DNS name or IP address of the manager you defined in Configure ASA FirePOWER Basic Settings, on page 107. For example, https://dc.example.com.

Alternatively, in ASDM, choose **Home** > **ASA FirePOWER Status** and click the link at the bottom of the dashboard.

For more information about ASA FirePOWER configuration, see the Management Center online help, *FireSIGHT System User Guide 5.4*, or *Firepower Management Center Configuration Guide 6.0* (available at http://www.cisco.com/c/en/us/support/security/defense-center/products-installation-and-configuration-guides-list.html).

ASDM

In ASDM, choose **Configuration** > **ASA FirePOWER Configuration**.

For more information about ASA FirePOWER configuration, see the module's online help in ASDM, *ASA FirePOWER Module User Guide 5.4*, or *ASA FirePOWER Services Local Management Configuration Guide 6.0* (available at http://www.cisco.com/c/en/us/support/security/asa-firepower-services/products-installation-and-configuration-guides-list.html.

Redirect Traffic to the ASA FirePOWER Module

For inline and inline tap (monitor-only) modes, you configure a service policy to redirect traffic to the module. If you want passive monitor-only mode, you configure a traffic redirection interface, which bypasses ASA policies.

The following topics explain how to configure these modes.

Configure Inline or Inline Tap Monitor-Only Modes

Redirect traffic to the ASA FirePOWER module by creating a service policy that identifies specific traffic that you want to send. In this mode, ASA policies, such as access rules, are applied to the traffic before it is redirected to the module.

Before you begin

- Be sure to configure consistent policies on the ASA and the ASA FirePOWER module. Both policies should reflect the inline or inline tap mode of the traffic.
- In multiple context mode, perform this procedure within each security context.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy Rules**.

Step 2 Choose **Add** > **Add Service Policy Rule**.

Step 3 Choose whether to apply the policy to a particular interface or apply it globally and click **Next**.

Step 4 Configure the traffic match. For example, you could match **Any Traffic** so that all traffic that passes your inbound access rules is redirected to the module. Or, you could define stricter criteria based on ports, ACL (source and destination criteria), or an existing traffic class. The other options are less useful for this policy. After you complete the traffic class definition, click **Next**.

Step 5 On the Rule Actions page, click the **ASA FirePOWER Inspection** tab.

Step 6 Check the **Enable ASA FirePOWER for this traffic flow** check box.

Step 7 In the If ASA FirePOWER Card Fails area, click one of the following:

- **Permit traffic**—Sets the ASA to allow all traffic through, uninspected, if the module is unavailable.
- **Close traffic**—Sets the ASA to block all traffic if the module is unavailable.

Step 8 (Optional) Check **Monitor-only** to send a read-only copy of traffic to the module (inline tap mode).

By default, the traffic is sent in inline mode. Be sure to configure consistent policies on the ASA and the ASA FirePOWER. Both policies should reflect the inline or monitor-only of the traffic.

Step 9 Click **Finish** and then **Apply**.

Repeat this procedure to configure additional traffic flows as desired.

Configure Passive Traffic Forwarding

If you want to operate the module in passive monitor-only mode, where the module gets a copy of the traffic and neither it nor the ASA can affect the network, configure a traffic forwarding interface and connect the interface to a SPAN port on a switch. For more details, see ASA FirePOWER Passive Monitor-Only Traffic Forwarding Mode, on page 99.

The following guidelines explain the requirements for this deployment mode:

- The ASA must be in single-context and transparent mode.
- You can configure up to 4 interfaces as traffic-forwarding interfaces. Other ASA interfaces can be used as normal.
- Traffic-forwarding interfaces must be physical interfaces, not VLANs or BVIs. The physical interface also cannot have any VLANs associated with it.

- Traffic-forwarding interfaces cannot be used for ASA traffic; you cannot name them or configure them for ASA features, including failover or management-only.
- You cannot configure both a traffic-forwarding interface and a service policy for ASA FirePOWER traffic.

Procedure

Step 1 Enter interface configuration mode for the physical interface you want to use for traffic-forwarding.

interface physical_interface

Example:

```
hostname(config)# interface gigabitethernet 0/5
```

Step 2 Remove any name configured for the interface. If this interface was used in any ASA configuration, that configuration is removed. You cannot configure traffic-forwarding on a named interface.

no nameif

Step 3 Enable traffic-forwarding.

traffic-forward sfr monitor-only

Note
You can ignore any warnings about traffic forwarding being for demonstration purposes only. This is a supported production mode.

Step 4 Enable the interface.

no shutdown

Repeat for any additional interfaces.

Example

The following example makes GigabitEthernet 0/5 a traffic-forwarding interface:

```
interface gigabitethernet 0/5
  no nameif
  traffic-forward sfr monitor-only
  no shutdown
```

Enable Captive Portal for Active Authentication

ASA FirePOWER includes identity policies that allow you to collect user identification information. By collecting user identity information, you can tailor access control rules to specific users and user groups,

selectively allowing and disallowing access based on the user. You can also analyze traffic based on user identity.

For HTTP/HTTPS connections, you can define identity rules that collect user identification through active authentication. If you want to implement active authentication identity rules, you must enable captive portal on the ASA to act as the authentication proxy port. When a connection matches an identity rule that requests active authentication, the ASA FirePOWER module redirects the authentication request to the ASA interface IP address/captive portal. The default port is 885, which you can change.

If you do not enable captive portal for the authentication proxy, only passive authentication is available.

Before you begin

- This feature is available in routed mode only for ASA FirePOWER 6.0+ only.
- In multiple context mode, perform this procedure within each security context.

Procedure

Step 1 Select **Tools** > **Command Line Tool**.

Step 2 Enable captive portal.

captive-portal {**global** | **interface** *name*} [**port** *number*]

Where:

- **global** enables captive portal globally on all interfaces.
- **interface** *name* enables captive portal on the specified interface only. You can enter the command multiple times to enable it on more than one interface. You can use this approach if you are redirecting traffic for only a subset of interfaces to the ASA FirePOWER module.
- **port** *number* optionally specifies the authentication port. If you do not include the keyword, port 885 is used. If you do include the keyword, the port number must be 1025 or higher.

Example:

For example, to enable captive portal globally on port 885, enter the following:

```
ciscoasa(config)# captive-portal global
ciscoasa(config)#
```

Step 3 In the ASA FirePOWER identity policy, ensure that the active authentication settings specify the same port you configured for captive portal, and configure the other required settings to enable active authentication.

Managing the ASA FirePOWER Module

This section includes procedures that help you manage the module.

Install or Reimage the Module

This section describes how to install or reimage a software module.

Install or Reimage the Software Module

If you purchase the ASA with the ASA FirePOWER module, the module software and required solid state drives (SSDs) come pre-installed and ready to configure. If you want to add the ASA FirePOWER software module to an existing ASA, or need to replace the SSD, you need to install the ASA FirePOWER boot software, partition the SSD, and install the system software according to this procedure.

Reimaging the module is the same procedure, except you should first uninstall the ASA FirePOWER module. You would reimage a system if you replace an SSD.

For information on how to physically install the SSD, see the ASA hardware guide.

Before you begin

- The free space on flash (disk0) should be at least 3GB plus the size of the boot software.
- In multiple context mode, perform this procedure in the system execution space.
- You must shut down any other software module that you might be running; the ASA can run a single software module at a time. You must do this from the ASA CLI.
- When reimaging the ASA FirePOWER module, use the **sw-module module shutdown** and **uninstall** commands to remove the old image. For example:

  sw-module module sfr shutdown
  sw-module module sfr uninstall
  reload

- If you have an active service policy redirecting traffic to an IPS or CX module, you must remove that policy. For example, if the policy is a global one, you could use **no service-policy** *ips_policy* **global**. If the service policy includes other rules you want to maintain, simply remove the redirection command from the relevant policy map, or the entire traffic class if redirection is the only action for the class. You can remove the policies using CLI or ASDM.
- If you have an active service policy redirecting traffic to another module, you must remove that policy. For example, if the policy is a global one, you could use **no service-policy** *module_policy* **global**.
- Obtain both the ASA FirePOWER Boot Image and System Software packages from Cisco.com.

Procedure

Step 1 Download the boot image to the ASA. Do not transfer the system software; it is downloaded later to the SSD. You have the following options:

- ASDM—First, download the boot image to your workstation, or place it on an FTP, TFTP, HTTP, HTTPS, SMB, or SCP server. Then, in ASDM, choose **Tools > File Management**, and then choose the appropriate **File Transfer** command, either **Between Local PC and Flash** or **Between Remote Server and Flash**. Transfer the boot software to disk0 on the ASA.

- ASA CLI—First, place the boot image on a TFTP, FTP, HTTP, or HTTPS server, then use the **copy** command to download it to flash. The following example uses TFTP.

```
ciscoasa# copy tftp://10.1.1.89/asasfr-5500x-boot-5.4.1-58.img
disk0:/asasfr-5500x-boot-5.4.1-58.img
```

Step 2 Download the ASA FirePOWER system software from Cisco.com to an HTTP, HTTPS, or FTP server accessible from the ASA FirePOWER management interface. Do not download it to disk0 on the ASA.

Step 3 Set the ASA FirePOWER module boot image location in ASA disk0 by entering the following command:

sw-module module sfr recover configure image disk0: *file_path*

Example:

```
hostname# sw-module module sfr recover configure image disk0:asasfr-5500x-boot-5.4.1-58.img
```

If you see a message like "ERROR: Another service (cxsc) is running, only one service is allowed to run at any time," it means that you already have a different software module configured. You must shut it down and remove it to install a new module as described in the prerequisites section above.

Step 4 Load the ASA FirePOWER boot image:

sw-module module sfr recover boot

Step 5 Wait approximately 5-15 minutes for the ASA FirePOWER module to boot up, and then open a console session to the now-running ASA FirePOWER boot image. You might need to press enter after opening the session to get to the login prompt. The default username is **admin** and the default password is **Admin123**.

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

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

If the module boot has not completed, the **session** command will fail with a message about not being able to connect over ttyS1. Wait and try again.

Step 6 Configure the system so that you can install the system software package:

asasfr-boot> **setup**

Example:

```
asasfr-boot> setup


                        Welcome to SFR Setup
                          [hit Ctrl-C to abort]
                      Default values are inside []
```

You are prompted for the following. Note that the management address and gateway, and DNS information, are the key settings to configure.

- Host name—Up to 65 alphanumeric characters, no spaces. Hyphens are allowed.
- Network address—You can set static IPv4 or IPv6 addresses, or use DHCP (for IPv4) or IPv6 stateless autoconfiguration.
- DNS information—You must identify at least one DNS server, and you can also set the domain name and search domain.
- NTP information—You can enable NTP and configure the NTP servers, for setting system time.

Step 7 Install the System Software image:

asasfr-boot> **system install** [**noconfirm**] *url*

Include the **noconfirm** option if you do not want to respond to confirmation messages. Use an HTTP, HTTPS, or FTP URL; if a username and password are required, you will be prompted to supply them.

When installation is complete, the system reboots. The time required for application component installation and for the ASA FirePOWER services to start differs substantially: high-end platforms can take 10 or more minutes, but low-end platforms can take 60-80 minutes or longer. (The **show module sfr** output should show all processes as Up.)

For example:

```
asasfr-boot> system install http://upgrades.example.com/packages/asasfr-sys-5.4.1-58.pkg
Verifying
Downloading
Extracting
Package Detail
        Description:                    Cisco ASA-FirePOWER 5.4.1-58 System Install
        Requires reboot:                Yes

Do you want to continue with upgrade? [y]: y
Warning: Please do not interrupt the process or turn off the system.
Doing so might leave system in unusable state.

Upgrading
Starting upgrade process ...
Populating new system image

Reboot is required to complete the upgrade. Press 'Enter' to reboot the system.
(press Enter)
Broadcast message from root (ttyS1) (Mon Feb 17 19:28:38 2014):

The system is going down for reboot NOW!
Console session with module sfr terminated.
```

Step 8 Open a session to the ASA FirePOWER module. You will see a different login prompt because you are logging into the fully functional module.

ciscoasa# **session sfr console**

Example:

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

Sourcefire ASA5555 v5.4.1 (build 58)
```

```
Sourcefire3D login:
```

Step 9 See Configure ASA FirePOWER Basic Settings, on page 107 to complete the setup.

Reset the Password

If you forget the password for the admin user, another user with CLI Configuration permissions can log in and change the password.

If there are no other users with the required permissions, you can reset the admin password from the ASA. The default password differs based on software release; see Defaults for ASA FirePOWER, on page 103.

Before you begin

- In multiple context mode, perform this procedure in the system execution space.
- The password-reset option on the ASA hw-module and sw-module commands does not work with ASA FirePOWER.

Procedure

Reset the module password for the user **admin** to the default:

session {1 | sfr} do password-reset

Use **1** for a hardware module, **sfr** for a software module.

Reload or Reset the Module

You can reload, or to reset and then reload, the module from the ASA.

Before you begin

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

Procedure

Enter the following command:

- **sw-module module sfr {reload | reset}**

Shut Down the Module

Shutting down the module software prepares the module to be safely powered off without losing configuration data.

Before you begin

- In multiple context mode, perform this procedure in the system execution space.
- If you reload the ASA, the module is not automatically shut down, so we recommend shutting down the module before reloading the ASA.

Procedure

Enter the following command:

- **sw-module module sfr shutdown**

Uninstall a Software Module Image

You can uninstall a software module image and its associated configuration.

Before you begin

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

Procedure

Step 1 Uninstall the software module image and associated configuration.

sw-module module sfr uninstall

Example:

```
ciscoasa# sw-module module sfr uninstall

Module sfr will be uninstalled. This will completely remove the disk image
associated with the sw-module including any configuration that existed within it.

Uninstall module sfr? [confirm]
```

Step 2 Reload the ASA.

reload

You must reload the ASA before you can install a new module.

Session to the Software Module From the ASA

Use the ASA FirePOWER CLI to configure basic network settings and to troubleshoot the module.

To access the ASA FirePOWER software module CLI from the ASA, you can session from the ASA.

You can either session to the module (using Telnet) or create a virtual console session. A console session might be useful if the control plane is down and you cannot establish a Telnet session. In multiple context mode, session from the system execution space.

In either a Telnet or a Console session, you are prompted for a username and password. You can log in with any username configured on the ASA FirePOWER. Initially, the **admin** username is the only one configured (and it is always available). The initial default password differs based on the type of image (full image or boot image) and software release; see Defaults for ASA FirePOWER, on page 103.

- Telnet session:

  session sfr

  When in the ASA FirePOWER CLI, to exit back to the ASA CLI, enter any command that would log you out of the module, such as **logout** or **exit**, or press **Ctrl-Shift-6, x**.

- Console session:

  session sfr console

  The only way out of a console session is to press **Ctrl-Shift-6, x**. Logging out of the module leaves you at the module login prompt.

Note Do not use the **session sfr console** command in conjunction with a terminal server where **Ctrl-Shift-6, x** is the escape sequence to return to the terminal server prompt. **Ctrl-Shift-6, x** is also the sequence to escape the ASA FirePOWER console and return to the ASA prompt. Therefore, if you try to exit the ASA FirePOWER console in this situation, you instead exit all the way to the terminal server prompt. If you reconnect the terminal server to the ASA, the ASA FirePOWER console session is still active; you can never exit to the ASA prompt. You must use a direct serial connection to return the console to the ASA prompt. Use the **session sfr** command instead of the console command when facing this situation.

Upgrade the System Software

Before applying an upgrade, ensure that the ASA is running the minimum required release for the new version; you might need to upgrade the ASA prior to upgrading the module. For more information about applying upgrades, see the Management Center online help, *FireSIGHT System User Guide 5.4*, or *Firepower Management Center Configuration Guide 6.0*.

For ASDM management, you can apply upgrades to the system software and components using **Configuration** > **ASA FirePOWER Configuration** > **Updates**. Click **Help** on the Updates page for more information.

Monitoring the ASA FirePOWER Module

The following topics provide guidance on monitoring the module. For ASA FirePOWER-related syslog messages, see the syslog messages guide. ASA FirePOWER syslog messages start with message number 434001.

Use **Tools > Command Line Interface** to use monitoring commands.

Showing Module Status

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

When managing the module with ASDM, you can also use the **Home > ASA FirePOWER Dashboard** page to view summary information about the software running on the module, product updates, licensing, system load, disk usage, system time, and interface status.

Showing Module Statistics

Use the show service-policy sfr command to display statistics and status for each service policy that includes the **sfr** command. Use **clear service-policy** to clear the counters.

The following example shows the ASA FirePOWER service policy and the current statistics as well as the module status. In monitor-only mode, the input counters remain at zero.

```
ciscoasa# show service-policy sfr

Global policy:
  Service-policy: global_policy
    Class-map: my-sfr-class
      SFR: card status Up, mode fail-close
        packet input 2626422041, packet output 2626877967, drop 0, reset-drop 0, proxied 0
```

Analyzing Operational Behavior (ASDM Management)

When you manage the ASA FirePOWER module using ASDM, you can view operational information for the module using these pages:

- **Home** > **ASA FirePOWER Reporting**—The reporting page provides Top 10 dashboards for a wide variety of module statistics, such as web categories, users, sources, and destinations for the traffic passing through the module.
- **Monitoring** > **ASA FirePOWER Monitoring**—There are several pages for monitoring the module, including syslog, task status, module statistics, and a real-time event viewer.

Monitoring Module Connections

To show connections through the ASA FirePOWER module, enter one of the following commands:

- **show asp table classify domain sfr**

  Shows the NP rules created to send traffic to the ASA FirePOWER module.

- **show asp drop**

  Shows dropped packets. The drop types are explained below.

- **show conn**

  Shows if a connection is being forwarded to a module by displaying the 'X - inspected by service module' flag.

The **show asp drop** command can include the following drop reasons related to the ASA FirePOWER module.

Frame Drops:

- sfr-bad-tlv-received—This occurs when ASA receives a packet from FirePOWER without a Policy ID TLV. This TLV must be present in non-control packets if it does not have the Standby/Active bit set in the actions field.
- sfr-request—The frame was requested to be dropped by FirePOWER due a policy on FirePOWER whereby FirePOWER would set the actions to Deny Source, Deny Destination, or Deny Pkt. If the frame should not have been dropped, review the policies on the module that are denying the flow.
- sfr-fail-close—The packet is dropped because the card is not up and the policy configured was 'fail-close' (rather than 'fail-open' which allows packets through even if the card was down). Check card status and attempt to restart services or reboot it.
- sfr-fail—The FirePOWER configuration was removed for an existing flow and we are not able to process it through FirePOWER it will be dropped. This should be very unlikely.
- sfr-malformed-packet—The packet from FirePOWER contains an invalid header. For instance, the header length may not be correct.
- sfr-ha-request—This counter is incremented when the security appliance receives a FirePOWER HA request packet, but could not process it and the packet is dropped.
- sfr-invalid-encap—This counter is incremented when the security appliance receives a FirePOWER packet with invalid message header, and the packet is dropped.
- sfr-bad-handle-received—Received Bad flow handle in a packet from FirePOWER Module, thus dropping flow. This counter is incremented, flow and packet are dropped on ASA as the handle for FirePOWER flow has changed in flow duration.
- sfr-rx-monitor-only—This counter is incremented when the security appliance receives a FirePOWER packet when in monitor-only mode, and the packet is dropped.

Flow Drops:

- sfr-request—The FirePOWER requested to terminate the flow. The actions bit 0 is set.
- reset-by-sfr—The FirePOWER requested to terminate and reset the flow. The actions bit 1 is set.

- sfr-fail-close—The flow was terminated because the card is down and the configured policy was 'fail-close'.

Example

The following is sample output from the **show asp table classify domain sfr** command. Note that if you configure redirection globally, only the input table will show data, and the output table will be empty. If you configure redirection by interface, both tables should contain data.

```
ciscoasa# show asp table classify domain sfr

Input Table
in  id=0x2aaae04034f0, priority=71, domain=sfr, deny=false
        hits=0, user_data=0x2aaadfdedf40, cs_id=0x0, use_real_addr, flags=0x0, protocol=0
        src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any
        dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any, dscp=0x0
        input_ifc=management, output_ifc=any

Output Table:

L2 - Output Table:

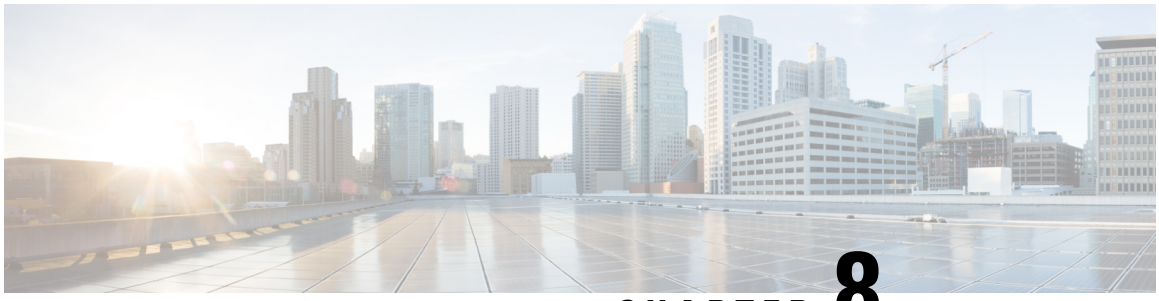
L2 - Input Table:

Last clearing of hits counters: Never
```

History for the ASA FirePOWER Module

Feature	Platform Releases	Description
ASA 5585-X (all models) support for the matching ASA FirePOWER SSP hardware module. ASA 5512-X through ASA 5555-X support for the ASA FirePOWER software module.	ASA 9.2(2.4) ASA FirePOWER 5.3.1	The ASA FirePOWER module supplies next-generation firewall services, including Next-Generation IPS (NGIPS), Application Visibility and Control (AVC), URL filtering, and Advanced Malware Protection (AMP).You can use the module in single or multiple context mode, and in routed or transparent mode. We introduced the following screens: Home > ASA FirePOWER Status Wizards > Startup Wizard > ASA FirePOWER Basic Configuration Configuration > Firewall > Service Policy Rules > Add Service Policy Rule > Rule Actions > ASA FirePOWER Inspection
ASA 5506-X support for the ASA FirePOWER software module, including support for configuring the module in ASDM	ASA 9.3(2) ASDM 7.3(3) ASA FirePOWER 5.4.1	You can run the ASA FirePOWER software module on the ASA 5506-X. You can manage the module using FireSIGHT Management Center, or you can use ASDM. We introduced the following screens: Home > ASA FirePOWER Dashboard, Home > ASA FirePOWER Reporting, Configuration > ASA FirePOWER Configuration (including sub-pages), Monitoring > ASA FirePOWER Monitoring (including sub-pages).

Feature	Platform Releases	Description
ASA FirePOWER passive monitor-only mode using traffic redirection interfaces	ASA 9.3(2) ASA FirePOWER 5.4.1	You can now configure a traffic forwarding interface to send traffic to the module instead of using a service policy. In this mode, neither the module nor the ASA affects the traffic. We fully supported the following command: **traffic-forward sfr monitor-only**. You can configure this in CLI only.
Support for managing the module through ASDM for the 5506H-X, 5506W-X, 5508-X, and 5516-X.	ASA 9.4(1) ASDM 7.4(1) ASA FirePOWER 5.4.1	You can manage the module using ASDM instead of using FireSIGHT Management Center. No new screens or commands were added.
Support for managing the module through ASDM for the 5512-X through 5585-X.	ASA 9.5.(1.5) ASDM 7.5(1.112) ASA FirePOWER 6.0	You can manage the module using ASDM instead of using Firepower Management Center (formerly FireSIGHT Management Center). No new screens or commands were added.
Captive portal for active authentication on ASA FirePOWER 6.0.	ASA 9.5.(2) ASA FirePOWER 6.0	The captive portal feature is required to enable active authentication using identity policies starting with ASA FirePOWER 6.0. We introduced or modified the following commands: **captive-portal**, **clear configure captive-portal**, **show running-config captive-portal**.
No support in 9.10(1) for the ASA FirePOWER module on the ASA 5506-X series and the ASA 5512-X	9.10(1)	The ASA 5506-X series and 5512-X no longer support the ASA FirePOWER module in 9.10(1) and later due to memory constraints. You must remain on 9.9(x) or lower to continue using this module. Other module types are still supported. If you upgrade to 9.10(1), the ASA configuration to send traffic to the FirePOWER module will be erased; make sure to back up your configuration before you upgrade. The FirePOWER image and its configuration remains intact on the SSD. If you want to downgrade, you can copy the ASA configuration from the backup to restore functionality.



CHAPTER 8

Cisco Umbrella

You can configure the device to redirect DNS requests to Cisco Umbrella, so that your FQDN policy defined in Cisco Umbrella can be applied to user connections. The following topics explain how to configure the Umbrella Connector to integrate the device with Cisco Umbrella.

- About Cisco Umbrella Connector, on page 125
- Licensing Requirements for Cisco Umbrella Connector, on page 126
- Guidelines and Limitations for Cisco Umbrella, on page 126
- Configure Cisco Umbrella Connector, on page 128
- Monitoring the Umbrella Connector , on page 133
- History for Cisco Umbrella Connector, on page 136

About Cisco Umbrella Connector

If you use Cisco Umbrella, you can configure the Cisco Umbrella Connector to redirect DNS queries to Cisco Umbrella. This allows Cisco Umbrella to identify requests to unpermitted or questionable domain names and apply your DNS-based security policy.

The Umbrella Connector is part of the system's DNS inspection. If your existing DNS inspection policy map decides to block or drop a request based on your DNS inspection settings, the request is not forwarded to Cisco Umbrella. Thus, you have two lines of protection: your local DNS inspection policy and your Cisco Umbrella cloud-based policy.

When redirecting DNS lookup requests to Cisco Umbrella, the Umbrella Connector adds an EDNS (Extension mechanisms for DNS) record. An EDNS record includes the device identifier information, organization ID, and client IP address. Your cloud-based policy can use those criteria to control access in addition to the reputation of the FQDN. You can also elect to encrypt the DNS request using DNSCrypt to ensure the privacy of usernames and internal IP addresses.

Cisco Umbrella Enterprise Security Policy

In your cloud-based Cisco Umbrella Enterprise Security policy, you can control access based on the reputation of the fully-qualified domain name (FQDN) in the DNS lookup request. Your Enterprise Security policy can enforce one of the following actions:

- Allow—If you have no block rules for an FQDN, and Cisco Umbrella determines that it belongs to a non-malicious site, then the site's actual IP address is returned. This is normal DNS lookup behavior.

- Proxy—If you have no block rules for an FQDN, and Cisco Umbrella determines that it belongs to a suspicious site, then the DNS reply returns the IP address of the Umbrella intelligent proxy. The proxy can then inspect the HTTP connection and apply URL filtering. You must ensure that intelligent proxy is enabled from the Cisco Umbrella dashboard (**Security Setting** > **Enable Intelligent Proxy**).
- Block—If you explicitly block an FQDN, or Cisco Umbrella determines that it belongs to a malicious site, then the DNS reply returns the IP address of the Umbrella cloud landing page for blocked connections.

Cisco Umbrella Registration

When you configure the Umbrella Connector on a device, it registers with Cisco Umbrella in the cloud. The registration process assigns a single device ID, which identifies one of the following:

- One standalone device in single context mode.
- One high availability pair in single context mode.
- One cluster in single context mode.
- One security context in a multiple-context standalone device.
- One security context of a high availability pair
- One security context of a cluster.

Once registered, the device details will appear on the Cisco Umbrella dashboard. You can then change which policy is attached to a device. During registration, either the policy you specify in the configuration is used, or the default policy is assigned. You can assign the same Umbrella policy to multiple devices. If you specify the policy, the device ID you receive differs from what you would get if you did not specify a policy.

Licensing Requirements for Cisco Umbrella Connector

To use the Cisco Umbrella Connector, you must have a 3DES license. If you are using Smart Licensing, your account must be enabled for export-controlled functionality.

The Cisco Umbrella portal has separate licensing requirements.

Guidelines and Limitations for Cisco Umbrella

Context Mode

- In multiple-context mode, you configure the Umbrella Connector in each context. Each context has a separate device ID, and is represented as a separate device in the Cisco Umbrella Connector dashboard. The device name is the hostname configured in the context, plus the hardware model, plus the context name. For example, CiscoASA-ASA5515-Context1.

Failover

- The active unit in the high availability pair registers the pair as a single unit with Cisco Umbrella. Both peers use the same device ID, which is formed from their serial numbers:

primary-serial-number_secondary-serial-number. For multiple context mode, each pair of security contexts is considered a single unit. You must configure high availability, and the units must have successfully formed a high-availability group (even if the standby device is currently in a failed state), before enabling Cisco Umbrella, or the registration will fail.

Cluster

- The cluster control unit registers the cluster as a single unit with Cisco Umbrella. All peers use the same device ID. For multiple context mode, a security context in the cluster is considered a single unit across all peers.

Additional Guidelines

- Redirection to Cisco Umbrella is done for DNS requests in through traffic only. DNS requests that the system itself initiates are never redirected to Cisco Umbrella. For example, FQDN-based access control rules are never resolved based on Umbrella policy, nor are any FQDNs that are used in other commands or configuration settings.
- The Cisco Umbrella Connector works on any DNS request in through traffic. However, the block and proxy actions are effective only if the DNS response is then used for HTTP/HTTPS connections, because the IP address returned is for a web site. Any blocked or proxied addresses for non-HTTP/HTTPS connections will either fail or complete in a misleading fashion. For example, pinging a blocked FQDN would result in pinging the server that hosts the Cisco Umbrella cloud block page.



Note Cisco Umbrella does try to intelligently identify FQDNs that might be non-HTTP/HTTPS, and does not return the IP address to the intelligent proxy for those FQDNs for proxied domain names.

- The system sends DNS/UDP traffic only to Cisco Umbrella. If you enable DNS/TCP inspection, the system does not send any DNS/TCP requests to Cisco Umbrella. However, DNS/TCP requests do not increment the Umbrella bypass counter.
- If you enable DNScrypt for Umbrella inspection, the system uses UDP/443 for the encrypted session. You must include UDP/443 along with UDP/53 in the class map that applies DNS inspection for Cisco Umbrella for DNScrypt to work correctly. Both UDP/443 and UDP/53 are included in the default inspection class for DNS, but if you create a custom class, ensure that you define an ACL that includes both ports for the match class.
- DNScrypt uses IPv4 only for the certificate update handshake. However, DNSscrypt does encrypt both IPv4 and IPv6 traffic.
- Cisco Umbrella and ASA FirePOWER processing are not compatible for a given connection. If you want to use both services, you must exclude UDP/53 and UDP/443 from ASA FirePOWER processing. For example, if you are currently redirecting all traffic to the ASA FirePOWER module, you must update the class to use an access list match. The access list should start by denying any connections to the Umbrella servers on destination ports UDP/53 and UDP/443, then permit any source to any destination. The ACL and match statements would be similar to the following:

```
access-list sfr extended deny udp any host 208.67.220.220 eq domain
access-list sfr extended deny udp any host 208.67.220.220 eq 443
```

```
access-list sfr extended permit ip any any

class-map sfr
 match access-list sfr
policy-map global_policy
 class sfr
  sfr fail-open
```

- There must be an IPv4 route to the Internet that can reach api.umbrella.com and api.opendns.com (registration uses IPv4 only). You also must have routes to the following DNS resolvers, and your access rules must allow DNS traffic to these hosts. These routes can go through either the data interfaces or the management interface; any valid route will work for both registration and DNS resolution. The default servers that the system uses are indicated; you can use the other servers by configuring the resolver in the Umbrella global settings.
  - 208.67.220.220 (system default for IPv4)
  - 208.67.222.222
  - 2620:119:53::53 (system default for IPv6)
  - 2620:119:35::35
- The system does not support the Umbrella FamilyShield service. If you configure the FamilyShield resolvers, you might get unexpected results.
- When evaluating whether to fail open, the system considers whether the Umbrella resolver is down, or if an intervening device drops the DNS request or response based on how long it has waited for the response after sending out the request. Other factors, such as no route to the Umbrella resolver, are not considered.
- To unregister a device, first delete the Umbrella configuration, then delete the device from the Cisco Umbrella dashboard.
- Any web requests that use IP addresses instead of FQDN will bypass Cisco Umbrella. In addition, if a roaming client obtains DNS resolution from a different WAN connection than the one that goes through an Umbrella-enabled device, connections that use those resolutions bypass Cisco Umbrella.
- If a user has an HTTP proxy, then the proxy might be doing DNS resolution, and the resolutions will not go through Cisco Umbrella.
- NAT DNS46 and DNS64 are not supported. You cannot translate DNS requests between IPv4 and IPv6 addressing.
- The EDNS record will include both the IPv4 and IPv6 host addresses.
- If the client uses DNS over HTTPS, then the cloud security service will not inspect DNS and HTTP/HTTPS traffic.

Configure Cisco Umbrella Connector

You can configure the device to interact with Cisco Umbrella in the cloud. The system redirects DNS lookup requests to Cisco Umbrella, which then applies your cloud-based Enterprise Security fully-qualified domain name (FQDN) policy. For malicious or suspicious traffic, users can be blocked from a site, or redirected to an intelligent proxy that can perform URL filtering based on your cloud-based policy.

The following procedure explains the end-to-end process for configuring the Cisco Umbrella Connector.

Before you begin

In multiple-context mode, perform this procedure in each security context that should use Cisco Umbrella.

Procedure

Step 1 Establish an account on Cisco Umbrella, https://umbrella.cisco.com.

Step 2 Install the CA Certificate from the Cisco Umbrella Registration Server, on page 129.

The device registration uses HTTPS, which requires that you install the root certificate.

Step 3 If it is not already enabled, configure DNS servers and enable DNS lookup on the interfaces.

Configure the settings on the **Configuration** > **Device Management** > **DNS** > **DNS Client** page.

You can use your own servers, or configure the Cisco Umbrella servers. DNS inspection automatically redirects to the Cisco Umbrella resolvers even if you configure different servers.

- 208.67.220.220
- 208.67.222.222
- 2620:119:53::53
- 2620:119:35::35

Step 4 Configure the Umbrella Connector Global Settings, on page 130.

Step 5 Enable Umbrella in the DNS Inspection Policy Map, on page 131.

Step 6 Verify the Umbrella Registration, on page 132.

Install the CA Certificate from the Cisco Umbrella Registration Server

You must import the root certificate to establish the HTTPS connection with the Cisco Umbrella registration server. The system uses the HTTPS connection when registering the device. In Cisco Umbrella choose **Deployments > Configuration > Root Certificate** and download the certificate.

Before you begin

When Umbrella updates its certificate, you need to download the new certificate. The root certificate might also change. Ensure that you have the correct root certificate uploaded.

When you update the certificate, you must disable Umbrella, then enable it again, so the system picks up the new certificate and registers correctly with Umbrella.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Advanced** > **Certificate Management** > **CA Certificates**.

Step 2 Click **Add**.

Step 3 Enter a **Trustpoint Name**, such as ctx1 or umbrella_server.

Step 4 Select **Paste Certificate in PEM Format**, then paste the certificate into the box.

It does not matter if you include the BEGIN CERTIFICATE and END CERTIFICATE lines.

Step 5 Click **Install Certificate**.

The certificate is created on the device. You will need to refresh your view to see the trustpoint listed.

Configure the Umbrella Connector Global Settings

The Umbrella global settings primarily define the API token that is needed to register the device with Cisco Umbrella.

The global settings are not sufficient to enable Umbrella. You must also enable Umbrella in your DNS inspection policy map, as described in Enable Umbrella in the DNS Inspection Policy Map, on page 131.

Before you begin

- Log into the Cisco Umbrella Network Devices Dashboard (https://login.umbrella.com/) and obtain a legacy network device API token for your organization. A token will be a hexadecimal string, for example, AABBA59A0BDE1485C912AFE. Generate a Legacy Network Devices API key from the Umbrella dashboard.
- Install the certificate for the Cisco Umbrella registration server.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Umbrella**.

Step 2 Select **Enable Umbrella**.

Step 3 Enter the API token in the **Token** field.

Step 4 (Optional.) If you intend to enable DNScrypt in the DNS inspection policy map, you can optionally configure the DNScrypt provider **Public Key** for certificate verification. If you do not configure the key, the default currently distributed public key is used for validation.

The key is a 32-byte hexadecimal value. Enter the hex value in ASCII with a colon separator for every 2 bytes. The key is 79 bytes long. Obtain this key from Cisco Umbrella.

The default key is:
B735:1140:206F:225D:3E2B:D822:D7FD:691E:A1C3:3CC8:D666:8D0C:BE04:BFAB:CA43:FB79

To revert to using the default public key, delete the key from the **Public Key** field.

Step 5 (Optional.) Select **EDNS Timeout** and change the idle timeout after which a connection from a client to the Umbrella server will be removed if there is no response from the server.

The timeout is in hours:minutes:seconds format, and can be from 0:0:0 to 1193:0:0. The default is 0:02:00 (2 minutes).

Step 6 (Optional.) In **Resolver IPv4** and **Resolver IPv6**, configure the addresses of the non-default Cisco Umbrella DNS servers, which resolve DNS requests, that you want to use.

If you do not configure these options, the system uses the default servers.

Step 7 (Optional.) Configure the local domain names for which Umbrella should be bypassed.

You can identify local domains for which DNS requests should bypass Cisco Umbrella and instead go directly to the configured DNS servers. For example, you can have your internal DNS server resolve all names for the organization's domain name on the assumption that all internal connections are allowed.

You can specify either a single regular expression class that includes the regular expression objects that define the local domains, or enter the names directly as regular expression objects. You can also combine these, although you can have at most one class.

Click the **Manage** button next to the **Local Domain Bypass Regex Class** option to create the class. You can also click the **Manage** button from the Add/Edit dialog box for regular expressions to create those objects.

Enable Umbrella in the DNS Inspection Policy Map

Configuring the global Umbrella settings is not enough to register the device and enable DNS lookup redirection. You must add Umbrella as part of your active DNS inspection.

You can enable Umbrella globally by adding it to the preset_dns_map DNS inspection policy map.

However, if you have customized DNS inspection and applied different inspection policy maps to different traffic classes, you must enable Umbrella on each class where you want the service.

The following procedure explains how to implement Umbrella globally. If you have customized DNS policy maps, please see Configure DNS Inspection Policy Map, on page 321.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **DNS**.

Step 2 Double-click the **preset_dns_map** inspection map to edit it.

Step 3 Click the **Umbrella Connections** tab and enable the connection to Cisco Umbrella in the cloud.

- **Umbrella**—Enables Cisco Umbrella. You can optionally specify the name of the Cisco Umbrella policy to apply to the device in the **Umbrella Tag** field. If you do not specify a policy, the default policy is applied. After registration, the Umbrella device ID is displayed next to the tag.
- **Enable Dnscrypt**—Enables DNScrypt to encrypt connections between the device and Cisco Umbrella. Enabling DNScrypt starts the key-exchange thread with the Umbrella resolver. The key-exchange thread performs the handshake with the resolver every hour and updates the device with a new secret key. Because DNScrypt uses UDP/443, you must ensure that the class map used for DNS inspection includes that port. Note that the default inspection class already includes UDP/443 for DNS inspection.
- **Fail Open**—Enable fail open if you want DNS resolution to work if the Umbrella DNS server is unavailable. When failing open, if the Cisco Umbrella DNS server is unavailable, Umbrella disables itself on this policy map and allows DNS requests to go to the other DNS servers configured on the system, if any. When the Umbrella DNS servers are available again, the policy map resumes using them.

If you do not select this option, DNS requests continue to go to the unreachable Umbrella resolver, so they will not get a response

Step 4 Click **OK**.

Verify the Umbrella Registration

After you configure the global Umbrella settings and enable Umbrella in DNS inspection, the device should contact Cisco Umbrella and register. You can check for successful registration by checking whether Cisco Umbrella provided a device ID.

Use **Tools** > **Command Line Interface** or an SSH session to enter these commands.

First, check the service policy statistics, and look for the Umbrella Registration line. This should indicate the policy applied by Cisco Umbrella (the tag), the HTTP status of the connection , and the device ID.

The status should be 200 Success. Error codes indicate the following: 401 indicates that the API token was incorrect, and 409 indicates that the device already exists in Cisco Umbrella.

Note that the Umbrella Resolver lines should not indicate that the resolvers are unresponsive. If they are, verify that you opened DNS communication to these IP addresses in your access control policy. This might be a temporary situation, or it might indicate a routing problem.

```
asa(config)# show service-policy inspect dns
Interface inside:
  Service-policy: global_policy
    Class-map: inspection_default
     Inspect: dns preset_dns_map, packet 0, lock fail 0, drop 0, reset-drop 0, 5-min-pkt-rate
 0 pkts/sec, v6-fail-close 0 sctp-drop-override 0
        message-length maximum client auto, drop 0
        message-length maximum 512, drop 0
        dns-guard, count 0
        protocol-enforcement, drop 0
        nat-rewrite, count 0
       umbrella registration: mode: fail-open tag: default, status: 200 success, device-id:
 010a13b8fbdfc9aa
          Umbrella ipv4 resolver: 208.67.220.220
          Umbrella ipv6 resolver: 2620:119:53::53
       Umbrella: bypass 0, req inject 0 - sent 0, res recv 0 - inject 0 local-domain-bypass
 10
        DNScrypt egress: rcvd 402, encrypt 402, bypass 0, inject 402
        DNScrypt ingress: rcvd 804, decrypt 402, bypass 402, inject 402
        DNScrypt: Certificate Update: completion 10, failure 1
```

You can also verify the running configuration (filter on policy-map). The umbrella command in the policy map updates to show the device ID. You cannot directly configure the device ID when you enable this command. The following example edits the output to show the relevant information. You can also see the device ID in ASDM by editing the DNS inspection map used for Umbrella; the ID is displayed on the **Umbrella Connections** tab.

```
ciscoasa(config)# show running-config policy-map
!
policy-map type inspect dns preset_dns_map
parameters
  message-length maximum client auto
```

```
  message-length maximum 512
  dnscrypt
  umbrella device-id 010a3e5760fdd6d3
  no tcp-inspection
policy-map global_policy
class inspection_default
  inspect dns preset_dns_map
```

Monitoring the Umbrella Connector

The following topics explain how to monitor the Umbrella Connector.

Monitoring the Umbrella Service Policy Statistics

You can view both summarized and detailed statistics for DNS inspection with Umbrella enabled.

Use **Tools** > **Command Line Interface** or an SSH session to enter these commands.

show service-policy inspect dns [**detail**]

Without the **detail** keyword, you see all the basic DNS inspection counters plus Umbrella configuration information. The status field provides the HTTP status code for the system's attempt to register with Cisco Umbrella.

The Resolver lines indicate which Umbrella servers are being used. These lines will say whether the server is **unresponsive**, or if the system is currently **probing** the server to determine if it has become available. If the mode is fail-open, the system allows DNS requests to go to other DNS servers (if configured); otherwise, DNS requests will not get a response so long as the Umbrella servers are unreponsive.

```
asa(config)# show service-policy inspect dns
Interface inside:
  Service-policy: global_policy
    Class-map: inspection_default
     Inspect: dns preset_dns_map, packet 0, lock fail 0, drop 0, reset-drop 0, 5-min-pkt-rate
 0 pkts/sec, v6-fail-close 0 sctp-drop-override 0
        message-length maximum client auto, drop 0
        message-length maximum 512, drop 0
        dns-guard, count 0
        protocol-enforcement, drop 0
        nat-rewrite, count 0
       umbrella registration: mode: fail-open tag: default, status: 200 success, device-id:
 010a13b8fbdfc9aa
          Umbrella ipv4 resolver: 208.67.220.220
          Umbrella ipv6 resolver: 2620:119:53::53
       Umbrella: bypass 0, req inject 0 - sent 0, res recv 0 - inject 0 local-domain-bypass
 10
        DNScrypt egress: rcvd 402, encrypt 402, bypass 0, inject 402
        DNScrypt ingress: rcvd 804, decrypt 402, bypass 402, inject 402
        DNScrypt: Certificate Update: completion 10, failure 1
```

The detailed output shows DNScrypt statistics and the keys used.

```
asa(config)# show service-policy inspect dns detail
Global policy:
  Service-policy: global_policy
    Class-map: inspection_default
```

```
   Class-map: dnscrypt30000
     Inspect: dns dns_umbrella, packet 12, lock fail 0, drop 0, reset-drop 0,
              5-min-pkt-rate 0 pkts/sec, v6-fail-close 0 sctp-drop-override 0
       message-length maximum client auto, drop 0
       message-length maximum 1500, drop 0
       dns-guard, count 3
       protocol-enforcement, drop 0
       nat-rewrite, count 0
      Umbrella registration: mode: fail-open tag: default, status: 200 SUCCESS, device-id:
010af97abf89abc3, retry 0
         Umbrella ipv4 resolver: 208.67.220.220
         Umbrella ipv6 resolver: 2620:119:53::53
      Umbrella: bypass 0, req inject 6 - sent 6, res recv 6 - inject 6 local-domain-bypass
10
         Umbrella app-id fail, count 0
         Umbrella flow alloc fail, count 0
         Umbrella block alloc fail, count 0
         Umbrella client flow expired, count 0
         Umbrella server flow expired, count 0
         Umbrella request drop, count 0
         Umbrella response drop, count 0
       DNScrypt egress: rcvd 6, encrypt 6, bypass 0, inject 6
       DNScrypt ingress: rcvd 18, decrypt 6, bypass 12, inject 6
         DNScrypt length error, count 0
         DNScrypt add padding error, count 0
         DNScrypt encryption error, count 0
         DNScrypt magic_mismatch error, count 0
         DNScrypt disabled, count 0
         DNScrypt flow error, count 0
         DNScrypt nonce error, count 0
       DNScrypt: Certificate Update: completion 1, failure 1
         DNScrypt Receive internal drop count 0
         DNScrypt Receive on wrong channel drop count 0
         DNScrypt Receive cannot queue drop count 0
         DNScrypt No memory to create channel count 0
         DNScrypt Send no output interface count 1
         DNScrypt Send open channel failed count 0
         DNScrypt Send no handle count 0
         DNScrypt Send dupb failure count 0
         DNScrypt Create cert update no memory count 0
         DNScrypt Store cert no memory count 0
         DNScrypt Certificate invalid length count 0
         DNScrypt Certificate invalid magic count 0
         DNScrypt Certificate invalid major version count 0
         DNScrypt Certificate invalid minor version count 0
         DNScrypt Certificate invalid signature count 0
         Last Successful: 01:42:29 UTC May 2 2018, Last Failed: None
         Magic DNSC, Major Version 0x0001, Minor Version 0x0000,
         Query Magic 0x714e7a696d657555, Serial Number 1517943461,
         Start Time 1517943461 (18:57:41 UTC Feb 6 2018)
         End Time 1549479461 (18:57:41 UTC Feb 6 2019)
         Server Public Key
240B:11B7:AD02:FAC0:6285:1E88:6EAA:44E7:AE5B:AD2F:921F:9577:514D:E226:D552:6836
         Client Secret Key Hash
48DD:E6D3:C058:D063:1098:C6B4:BA6F:D8A7:F0F8:0754:40B0:AFB3:CB31:2B22:A7A4:9CEE
         Client Public key
6CB9:FA4B:4273:E10A:8A67:BA66:76A3:BFF5:2FB9:5004:CD3B:B3F2:86C1:A7EC:A0B6:1A58
         NM key Hash
9182:9F42:6C01:003C:9939:7741:1734:D199:22DF:511E:E8C9:206B:D0A3:8181:CE57:8020
```

Monitoring Umbrella Syslog Messages

You can monitor the following Umbrella-related syslog messages:

- %ASA-3-339001: DNSCRYPT certificate update failed for *number* tries.

  Check that there is a route to the Umbrella server and that the egress interface is up and functioning correctly. Also check that the public key configured for DNScrypt is correct. You might need to obtain a new key from Cisco Umbrella.

- %ASA-3-339002: Umbrella device registration failed with error code *error_code*.

  The error codes have the following meanings:

  - 400—There is a problem with the request format or content. The token is probably too short or corrupted. Verify that the token matches the one on the Umbrella Dashboard.
  - 401—The API token is not authorized. Try reconfiguring the token. If you refreshed the token on the Umbrella Dashboard, then you must ensure that you use the new token.
  - 409—The device ID conflicts with another organization. Please check with the Umbrella Administrator to see what the issue might be.
  - 500—There is an internal server error. Check with the Umbrella Administrator to see what the issue might be.

- %ASA-6-339003: Umbrella device registration was successful.
- %ASA-3-339004: Umbrella device registration failed due to missing token.

  You must obtain an API token from Cisco Umbrella and configure it in the global Umbrella settings.

- %ASA-3-339005: Umbrella device registration failed after *number* retries.

  Check the syslog 339002 messages to identify the errors that you need to fix.

- %ASA-3-339006: Umbrella resolver *IP_address* is reachable, resuming Umbrella redirect.

  This message indicates that the system is functioning normally again. No action is needed.

- %ASA-3-339007: Umbrella resolver *IP_address* is unresponsive and fail-close mode used, starting probe to resolver.

  Because you are using fail-close mode, users will not get responses to their DNS requests until the Umbrella DNS server comes back online. If the problem persists, verify that there is a route from the system to the Umbrella servers, and that you allow DNS traffic to the servers in your access control policy.

History for Cisco Umbrella Connector

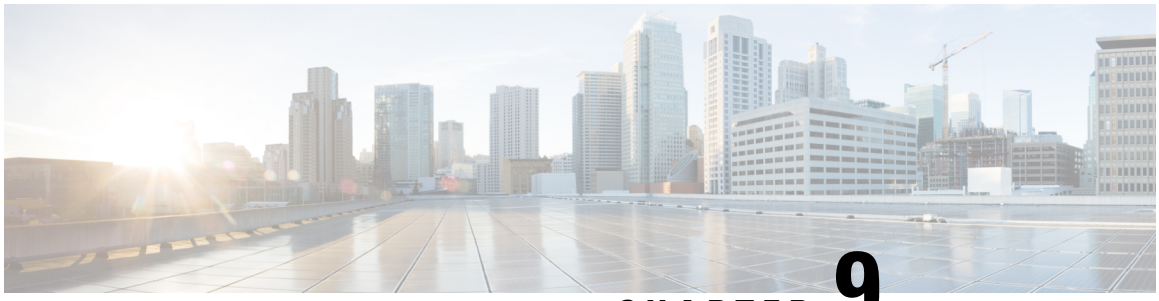
Feature Name	Platform Releases	Description
Cisco Umbrella support.	9.10(1)	You can configure the device to redirect DNS requests to Cisco Umbrella, so that your Enterprise Security policy defined in Cisco Umbrella can be applied to user connections. You can allow or block connections based on FQDN, or for suspicious FQDNs, you can redirect the user to the Cisco Umbrella intelligent proxy, which can perform URL filtering. The Umbrella configuration is part of the DNS inspection policy. We added or modified the following screens: **Configuration** > **Firewall** > **Objects** > **Umbrella**, **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **DNS**.
Cisco Umbrella Enhancements.	9.12(1)	You can now identify local domain names that should bypass Cisco Umbrella. DNS requests for these domains go directly to the DNS servers without Umbrella processing. You can also identify which Umbrella servers to use for resolving DNS requests. Finally, you can define the Umbrella inspection policy to fail open, so that DNS requests are not blocked if the Umbrella server is unavailable. We modified the following screens: **Configuration** > **Firewall** > **Objects** > **Umbrella**, **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **DNS**.



PART II

Firewall Services for Virtual Environments

- Attribute-Based Access Control, on page 139



CHAPTER 9

Attribute-Based Access Control

Attributes are customized network objects for use in your configuration. You can define and use them in ASA configurations to filter traffic associated with one or more virtual machines in an VMware ESXi environment managed by VMware vCenter. Attributes allow you to define access control lists (ACLs) to assign policies to traffic from groups of virtual machines sharing one or more attributes. You assign attributes to virtual machines within the ESXi environment and configure an attribute agent, which connects to vCenter or a single ESXi host using HTTPS. The agent then requests and retrieves one or more bindings which correlate specific attributes to the primary IP address of a virtual machine.

Attribute-based access control is supported on all hardware platforms, and on all ASAv platforms running on ESXi, KVM, or HyperV hypervisors. Attributes can only be retrieved from virtual machines running on an ESXi hypervisor.

- Guidelines for Attribute-Based Network Objects, on page 139
- Configure Attribute-Based Access Control, on page 140
- Monitoring Attribute-Based Network Objects , on page 145
- History for Attribute-Based Access Control, on page 145

Guidelines for Attribute-Based Network Objects

IPv6 Guidelines

- IPv6 addresses not supported by vCenter for host credentials.
- IPv6 is supported for virtual machine bindings where the primary IP address of the virtual machine is an IPv6 address.

Additional Guidelines and Limitations

- Multi-context mode is not supported. Attribute-based network objects are supported for single-mode context only.
- Attribute-based network objects support binding to the virtual machine's primary address only. Binding to multiple vNICs on a single virtual machine is not supported.
- Attribute-based network objects may only be configured for objects used for access groups. Network objects for other features (NAT, etc.) are not supported.

- Virtual machines must be running VMware Tools in order to report primary IP addresses to vCenter. The ASA is not notified of attribute changes unless vCenter knows the IP address of the virtual machine. This is a vCenter restriction.
- Attribute-based network objects are not supported in the Amazon Web Services (AWS) or Microsoft Azure public cloud environments.

Configure Attribute-Based Access Control

The following procedure provides a general sequence for implementing attribute-based access control on managed virtual machines in a VMware ESXi environment.

Procedure

Step 1 Assign custom attribute types and values to your managed virtual machines. See Configure Attributes for vCenter Virtual Machines, on page 140.

Step 2 Configure an attribute agent to connect to your vCenter Server or ESXi host. See Configure a VM Attribute Agent, on page 142.

Step 3 Configure attribute-based network objects needed for your deployment scheme. See Configure Attribute-Based Network Objects, on page 143.

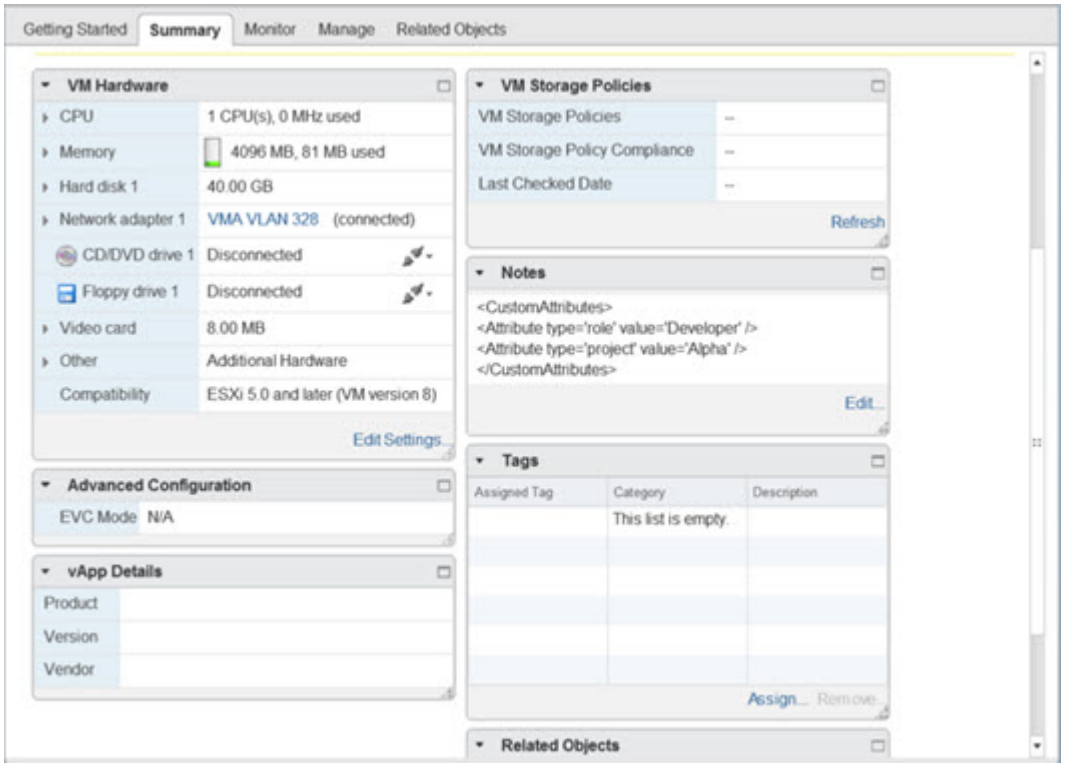
Step 4 Configure the access control lists and rules. See Configure Access Rules Using Attribute-Based Network Objects, on page 144 .

Configure Attributes for vCenter Virtual Machines

You assign custom attribute types and values to virtual machines, and associate these attributes to network objects. You can then use these attribute-based network objects to apply ACLs to a set of virtual machines with common user-defined characteristics. For example, you could isolate developer build machines from test machines, or group virtual machines by project and/or location. For the ASA to monitor virtual machines using attributes, you need to make the attributes available to vCenter from the managed virtual machines. You do this by inserting a formatted text file into the Notes field, which is found on the Summary page of virtual machines in vCenter.

You can see the Notes field in the following figure.

Figure 16: Summary Tab of a Virtual Machine in vCenter



To specify custom attributes, you copy a properly formatted XML file into the Notes field for the virtual machine. The format of the file is:

```
<CustomAttributes>
<Attribute type='attribute-type' value='attribute-value'/>
...
</CustomAttributes>
```

A single virtual machine may have multiple attributes defined by repeating the second line above. Note that each line must identify a unique attribute type. If the same attribute type is defined with multiple attribute values, each binding update for that attribute type will overwrite the previous one.

For string attribute values, the value associated with the object definition must be an exact match to the value reported to vCenter by the virtual machine. For example, an attribute value *Build Machine* does not match the annotation value *build machine* on the virtual machine. A binding would not be added to the host-map for this attribute.

You can define multiple unique attribute types in a single file.

Procedure

Step 1 Select the virtual machine from your vCenter inventory.

Step 2 Click the **Summary** tab for the virtual machine.

Step 3 In the **Notes** field, click the **Edit** link.

Step 4 Paste the custom attributes text file into the **Edit Notes** box. The text file should follow the XML template format:

Example:

```
<CustomAttributes>
<Attribute type='attribute-type' value='attribute-value'/>
...
</CustomAttributes>
```

Step 5 Click **OK**.

Configure a VM Attribute Agent

You configure a VM attribute agent to communicate with vCenter or a single ESXi host. When you assign attributes to virtual machines within the VMware environment, the attribute agent sends a message to vCenter indicating which attributes have been configured, and vCenter responds with a binding update for every virtual machine where a matching attribute type is configured.

The VM attribute agent and vCenter exchange binding updates as follows:

- If the agent issues a request containing a new attribute type, vCenter responds with a binding update for every virtual machine where the attribute type is configured. After that point, vCenter only issues a new binding when an attribute value is added or changed.
- If a monitored attribute changes for one or more virtual machines, a binding update message is received. Each binding message is identified by the IP address of the virtual machine reporting the attribute value.
- If multiple attributes are being monitored by a single agent, a single binding update contains the current value of all monitored attributes for each virtual machine.
- If a specific attribute being monitored by the agent is not configured on a virtual machine, the binding will contain an empty attribute value for that virtual machine.
- If a virtual machine has not been configured with any monitored attributes, vCenter does not send a binding update.

Each attribute agent communicates with exactly one vCenter or ESXi host. A single ASA may have multiple attribute agents defined, each communicating with a different vCenter, or one or more communicating with the same vCenter.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **VM Attribute Agent**.

Step 2 Click **Add**.

Step 3 In the Host Information area:

a) Choose whether to enable IP address and authentication credentials.
b) Enter a DNS host name or IP address.
c) Enter a user name.
d) Select the password type: **Clear Text**, **UnEncrypted**, or **Encrypted**.

e) Enter the password.

Step 4 In the Keepalive Information area:

a) Enter the **Retry Interval**. Enter a value between 1 and 65535. The default is 30.

b) Enter the **Retry Count**. Enter a value between 1 and 32. The default is 3.

Step 5 Click **OK**.

Configure Attribute-Based Network Objects

Attribute-based network objects filter traffic according to attributes associated with one or more virtual machines in a VMware ESXi environment. You can define access control lists (ACLs) to assign policies to traffic from groups of virtual machines sharing one or more attributes.

For example, you can configure access rules that permit machines with an *engineering* attribute to access machines with a *eng_lab* attribute. A network admininstrator can add or remove engineering machines and lab servers while the security policy managed by the security administrator continues to work automatically without manual updates to the access rules.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Access Rules** > **Advanced Options**.

Step 2 Check the **Enable Object Group Search Algorithm** check box.

You must enable object group search to configure VM attributes.

Step 3 Choose **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups**.

Step 4 Do one of the following:

- Choose **Add > Network Object Attributes** to add a new attribute-based network object. Enter a name and optionally, a description.
- Choose an existing attribute-based network object and click **Edit**.

Step 5 For a new attribute-based network object, enter values for the following fields:

a) **Agent Name**—Click the browse button and select a VM attribute agent (or define a new one).

If you configure a attribute-based network object to use an attribute agent which has not been configured, a placeholder agent is automatically created with no credentials and default keepalive values. This agent remains in the "No credentials available" state until host credentials are supplied.

b) **Attribute Type**—This string entry defines the attribute type and must include the **custom.** prefix. For example, custom.role.

c) **Attribute Value**—This string entry associates a value to the attribute type.

Together, the *Attribute Type* and *Attribute Value* pair define a unique attribute. This allows you to define multiple attributes that suit your particular deployment scheme. If you define the same attribute type more than once with multiple attribute values, the last value defined overwrites the previous one.

Step 6 Click **OK** .

Configure Access Rules Using Attribute-Based Network Objects

To apply an access rule using attribute-based network objects, perform the following steps.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Access Rules**.

The rules are organized by interface and direction, with a separate group for global rules. If you configure management access rules, they are repeated on this page. These groups are equivalent to the extended ACL that is created and assigned to the interface or globally as an access group. These ACLs also appear on the ACL Manager page.

Step 2 Do any of the following:

- To add a new rule, choose **Add > Add Access Rule**.
- To insert a rule at a specific location within a container, select an existing rule and choose **Add > Insert** to add the rule above it, or choose **Add > Insert After**.
- To edit a rule, select it and click **Edit**.

Step 3 Fill in the rule properties. The primary options to select are:

- **Interface**—The interface to which the rule applies. Select Any to create a global rule. For bridge groups in routed mode, you can create access rules for both the Bridge Virtual Interface (BVI) and each bridge group member interface.
- **Action: Permit/Deny**—Whether you are permitting (allowing) the described traffic or are denying (dropping) it.
- **Source/Destination criteria**—Select the source attribute-based network object (originating object) and destination attribute-based network object (target object of the traffic flow). You can also specify a user or user group name for the source. Additionally, you can use the Service field to identify the specific type of traffic if you want to focus the rule more narrowly than all IP traffic. If you implement Trustsec, you can use security groups to define source and destination.

For detailed information on all of the available options, see Access Rule Properties, on page 20.

When you are finished defining the rule, click **OK** to add the rule to the table.

Step 4 Click **Apply** to save the access rule to your configuration.

Monitoring Attribute-Based Network Objects

For attribute-based network objects, you can analyze the usage of an individual object. From their page in the **Configuration > Firewall > Objects> Network Objects/Groups** folder, click the **Where Used** button.

For attribute-based network objects, you can also click the Not Used button to find objects that are not used in any rules. This display gives you a short-cut for deleting these unused objects.

History for Attribute-Based Access Control

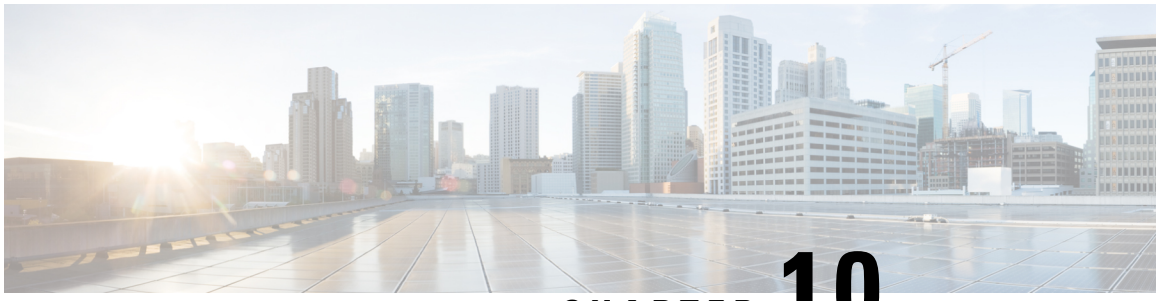
Feature Name	Platform Releases	Description
Support for Attribute-Based Network Objects	9.7.(1)	You can now control network access using virtual machine attributes in addition to traditional network charateristics such as IP addresses, protocols, and ports. The virtual machines must be in a VMware ESXi environment. We introduced or modified the following screens: **Configuration** > **Firewall** > **Objects** > **Network Object Attributes**. We introduced the following screen: **Configuration** > **Firewall** > **VM Attribute Agent**.
Remove support for VM attribute-based network objects from ASA 5506-X (all models), 5508-X, 5512-X, 5516-X.	9.10(1)	You can no longer use VM-attribute based network objects on the following platforms: ASA 5506-X (all models), 5508-X, 5512-X, 5516-X.



PART III

Network Address Translation

- Network Address Translation (NAT), on page 149
- NAT Examples and Reference, on page 215
- Mapping Address and Port (MAP), on page 275



CHAPTER 10

Network Address Translation (NAT)

The following topics explain Network Address Translation (NAT) and how to configure it.

- Why use NAT?, on page 149
- NAT basics, on page 150
- Guidelines for NAT, on page 155
- Dynamic NAT, on page 162
- Dynamic PAT, on page 170
- Static NAT, on page 189
- Identity NAT, on page 201
- Monitoring NAT, on page 208
- History for NAT, on page 208

Why use NAT?

Each computer and device within an IP network is assigned a unique IP address that identifies the host. Because of a shortage of public IPv4 addresses, most of these IP addresses are private, not routable anywhere outside of the private company network. RFC 1918 defines the private IP addresses you can use internally that should not be advertised:

- 10.0.0.0 through 10.255.255.255
- 172.16.0.0 through 172.31.255.255
- 192.168.0.0 through 192.168.255.255

One of the main functions of NAT is to enable private IP networks to connect to the Internet. NAT replaces a private IP address with a public IP address, translating the private addresses in the internal private network into legal, routable addresses that can be used on the public Internet. In this way, NAT conserves public addresses because it can be configured to advertise at a minimum only one public address for the entire network to the outside world.

Other functions of NAT include:

- Security—Keeping internal IP addresses hidden discourages direct attacks.
- IP routing solutions—Overlapping IP addresses are not a problem when you use NAT.

- Flexibility—You can change internal IP addressing schemes without affecting the public addresses available externally; for example, for a server accessible to the Internet, you can maintain a fixed IP address for Internet use, but internally, you can change the server address.
- Translating between IPv4 and IPv6 (Routed mode only)—If you want to connect an IPv6 network to an IPv4 network, NAT lets you translate between the two types of addresses.



Note NAT is not required. If you do not configure NAT for a given set of traffic, that traffic will not be translated, but will have all of the security policies applied as normal.

NAT basics

The following topics explain some of the basics of NAT.

NAT terminology

This document uses the following terminology:

- Real address/host/network/interface—The real address is the address that is defined on the host, before it is translated. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the inside network would be the "real" network. Note that you can translate any network connected to the device, not just an inside network. Therefore if you configure NAT to translate outside addresses, "real" can refer to the outside network when it accesses the inside network.
- Mapped address/host/network/interface—The mapped address is the address that the real address is translated to. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the outside network would be the "mapped" network.



Note During address translation, IP addresses configured for the device interfaces are not translated.

- Bidirectional initiation—Static NAT allows connections to be initiated *bidirectionally*, meaning both to the host and from the host.
- Source and destination NAT—For any given packet, both the source and destination IP addresses are compared to the NAT rules, and one or both can be translated/untranslated. For static NAT, the rule is bidirectional, so be aware that "source" and "destination" are used in commands and descriptions throughout this guide even though a given connection might originate at the "destination" address.

NAT types

You can implement NAT using the following methods:

- Dynamic NAT—A group of real IP addresses are mapped to a (usually smaller) group of mapped IP addresses, on a first come, first served basis. Only the real host can initiate traffic. See Dynamic NAT, on page 162.
- Dynamic Port Address Translation (PAT)—A group of real IP addresses are mapped to a single IP address using a unique source port of that IP address. See Dynamic PAT, on page 170.
- Static NAT—A consistent mapping between a real and mapped IP address. Allows bidirectional traffic initiation. See Static NAT, on page 189.
- Identity NAT—A real address is statically translated to itself, essentially bypassing NAT. You might want to configure NAT this way when you want to translate a large group of addresses, but then want to exempt a smaller subset of addresses. See Identity NAT, on page 201.

Network object NAT and Twice NAT

You can implement address translation in two ways: *network object NAT* and *twice NAT*.

We recommend using network object NAT unless you need the extra features that twice NAT provides. It is easier to configure network object NAT, and it might be more reliable for applications such as Voice over IP (VoIP). (For VoIP, you might see a failure in the translation of indirect addresses that do not belong to either of the objects used in the rule.)

Network object NAT

All NAT rules that are configured as a parameter of a network object are considered to be *network object NAT* rules. This is a quick and easy way to configure NAT for a network object. You cannot create these rules for a group object, however.

After you configure the network object, you can then identify the mapped address for that object, either as an inline address or as another network object or network object group.

When a packet enters an interface, both the source and destination IP addresses are checked against the network object NAT rules. The source and destination address in the packet can be translated by separate rules if separate matches are made. These rules are not tied to each other; different combinations of rules can be used depending on the traffic.

Because the rules are never paired, you cannot specify that sourceA/destinationA should have a different translation than sourceA/destinationB. Use twice NAT for that kind of functionality, where you can identify the source and destination address in a single rule.

Twice NAT

Twice NAT lets you identify both the source and destination address in a single rule. Specifying both the source and destination addresses lets you specify that sourceA/destinationA can have a different translation than sourceA/destinationB.

Note For static NAT, the rule is bidirectional, so be aware that "source" and "destination" are used in commands and descriptions throughout this guide even though a given connection might originate at the "destination" address. For example, if you configure static NAT with port address translation, and specify the source address as a Telnet server, and you want all traffic going to that Telnet server to have the port translated from 2323 to 23, then you must specify the *source* ports to be translated (real: 23, mapped: 2323). You specify the source ports because you specified the Telnet server address as the source address.

The destination address is optional. If you specify the destination address, you can either map it to itself (identity NAT), or you can map it to a different address. The destination mapping is always a static mapping.

Comparing Network object NAT and Twice NAT

The main differences between these two NAT types are:

- How you define the real address.
  - Network object NAT—You define NAT as a parameter for a network object. A network object names an IP host, range, or subnet so you can then use the object in the NAT configuration instead of the actual IP addresses. The network object IP address serves as the real address. This method lets you easily add NAT to network objects that might already be used in other parts of your configuration.
  - Twice NAT—You identify a network object or network object group for both the real and mapped addresses. In this case, NAT is not a parameter of the network object; the network object or group is a parameter of the NAT configuration. The ability to use a network object *group* for the real address means that twice NAT is more scalable.
- How source and destination NAT is implemented.
  - Network object NAT— Each rule can apply to either the source or destination of a packet. So two rules might be used, one for the source IP address, and one for the destination IP address. These two rules cannot be tied together to enforce a specific translation for a source/destination combination.
  - Twice NAT—A single rule translates both the source and destination. A packet matches one rule only, and further rules are not checked. Even if you do not configure the optional destination address, a matching packet still matches one twice NAT rule only. The source and destination are tied together, so you can enforce different translations depending on the source/destination combination. For example, sourceA/destinationA can have a different translation than sourceA/destinationB.
- Order of NAT Rules.
  - Network object NAT—Automatically ordered in the NAT table.
  - Twice NAT—Manually ordered in the NAT table (before or after network object NAT rules).

NAT rule order

Network object NAT and twice NAT rules are stored in a single table that is divided into three sections. Section 1 rules are applied first, then section 2, and finally section 3, until a match is found. For example, if

a match is found in section 1, sections 2 and 3 are not evaluated. The following table shows the order of rules within each section.



Note There is also a Section 0, which contains any NAT rules that the system creates for its own use. These rules have priority over all others. The system automatically creates these rules and clears xlates as needed. You cannot add, edit, or modify rules in Section 0.

Table 9: NAT Rule Table

Table Section	Rule Type	Order of Rules within the Section
Section 1	Twice NAT	Applied on a first match basis, in the order they appear in the configuration. Because the first match is applied, you must ensure that specific rules come before more general rules, or the specific rules might not be applied as desired. By default, twice NAT rules are added to section 1. By "specific rules first," we mean: • Static rules should come before dynamic rules. • Rules that include destination translation should come before rules with source translation only. If you cannot eliminate overlapping rules, where more than one rule might apply based on the source or destination address, be especially careful to follow these recommendations. If you do not follow the above recommendation, ensure that accurate routes are configured for the mapped addresses to match the correct NAT rule.
Section 2	Network object NAT	If a match in section 1 is not found, section 2 rules are applied in the following order: **1.** Static rules. **2.** Dynamic rules. **Within each rule type, the following ordering guidelines are used:** **1.** Quantity of real IP addresses—From smallest to largest. For example, an object with one address will be assessed before an object with 10 addresses. **2.** For quantities that are the same, then the IP address number is used, from lowest to highest. For example, 10.1.1.0 is assessed before 11.1.1.0. **3.** If the same IP address is used, then the name of the network object is used, in alphabetical order. For example, abracadabra is assessed before catwoman.

Table Section	Rule Type	Order of Rules within the Section
Section 3	Twice NAT	If a match is still not found, section 3 rules are applied on a first match basis, in the order they appear in the configuration. This section should contain your most general rules. You must also ensure that any specific rules in this section come before general rules that would otherwise apply.

For section 2 rules, for example, you have the following IP addresses defined within network objects:

- 192.168.1.0/24 (static)
- 192.168.1.0/24 (dynamic)
- 10.1.1.0/24 (static)
- 192.168.1.1/32 (static)
- 172.16.1.0/24 (dynamic) (object def)
- 172.16.1.0/24 (dynamic) (object abc)

The resultant ordering would be:

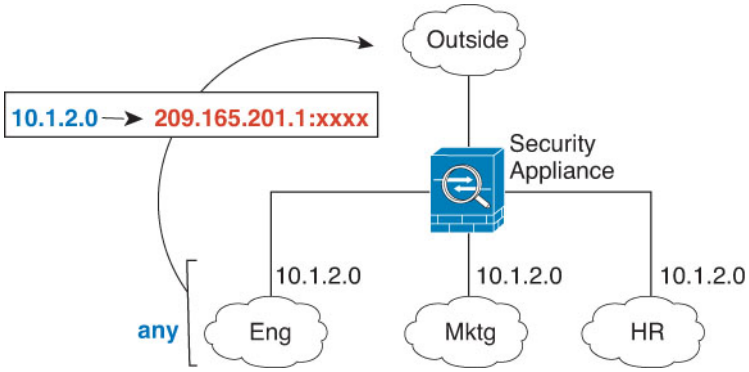
- 192.168.1.1/32 (static)
- 10.1.1.0/24 (static)
- 192.168.1.0/24 (static)
- 172.16.1.0/24 (dynamic) (object abc)
- 172.16.1.0/24 (dynamic) (object def)
- 192.168.1.0/24 (dynamic)

NAT interfaces

Except for bridge group member interfaces, you can configure a NAT rule to apply to any interface (in other words, all interfaces), or you can identify specific real and mapped interfaces. You can also specify any interface for the real address, and a specific interface for the mapped address, or vice versa.

For example, you might want to specify any interface for the real address and specify the outside interface for the mapped address if you use the same private addresses on multiple interfaces, and you want to translate them all to the same global pool when accessing the outside.

Figure 17: Specifying Any Interface



However, the concept of "any" interface does not apply to bridge group member interfaces. When you specify "any" interface, all bridge group member interfaces are excluded. Thus, to apply NAT to bridge group members, you must specify the member interface. This could result in many similar rules where only one interface is different. You cannot configure NAT for the Bridge Virtual Interface (BVI) itself, you can configure NAT for member interfaces only.

Guidelines for NAT

The following topics provide detailed guidelines for implementing NAT.

Firewall mode guidelines for NAT

NAT is supported in routed and transparent firewall mode.

However, configuring NAT on bridge group member interfaces (interfaces that are part of a Bridge Group Virtual Interface, or BVI) has the following restrictions:

- When configuring NAT for the members of a bridge group, you specify the member interface. You cannot configure NAT for the bridge group interface (BVI) itself.
- When doing NAT between bridge group member interfaces, you must specify the real and mapped addresses. You cannot specify "any" as the interface.
- You cannot configure interface PAT when the mapped address is a bridge group member interface, because there is no IP address attached to the interface.
- You cannot translate between IPv4 and IPv6 networks (NAT64/46) when the source and destination interfaces are members of the same bridge group. Static NAT/PAT 44/66, dynamic NAT44/66, and dynamic PAT44 are the only allowed methods; dynamic PAT66 is not supported. However, you can do NAT64/46 between members of different bridge groups, or between a bridge group member (source) and standard routed interface (destination).

IPv6 NAT guidelines

NAT supports IPv6 with the following guidelines and restrictions.

- For standard routed mode interfaces, you can also translate between IPv4 and IPv6.

- You cannot translate between IPv4 and IPv6 for interfaces that are members of the same bridge group. You can translate between two IPv6 or two IPv4 networks only. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.
- You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.
- For static NAT, you can specify an IPv6 subnet up to /64. Larger subnets are not supported.
- When using FTP with NAT46, when an IPv4 FTP client connects to an IPv6 FTP server, the client must use either the extended passive mode (EPSV) or extended port mode (EPRT); PASV and PORT commands are not supported with IPv6.

IPv6 NAT best practices

You can use NAT to translate between IPv6 networks, and also to translate between IPv4 and IPv6 networks (routed mode only). We recommend the following best practices:

- NAT66 (IPv6-to-IPv6)—We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (twice NAT only).
- NAT46 (IPv4-to-IPv6)—We recommend using static NAT. Because the IPv6 address space is so much larger than the IPv4 address space, you can easily accommodate a static translation. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (twice NAT only). When translating to an IPv6 subnet (/96 or lower), the resulting mapped address is by default an IPv4-embedded IPv6 address, where the 32-bits of the IPv4 address is embedded after the IPv6 prefix. For example, if the IPv6 prefix is a /96 prefix, then the IPv4 address is appended in the last 32-bits of the address. For example, if you map 192.168.1.0/24 to 201b::0/96, then 192.168.1.4 will be mapped to 201b::0.192.168.1.4 (shown with mixed notation). If the prefix is smaller, such as /64, then the IPv4 address is appended after the prefix, and a suffix of 0s is appended after the IPv4 address. You can also optionally translate the addresses net-to-net, where the first IPv4 address maps to the first IPv6 address, the second to the second, and so on.
- NAT64 (IPv6-to-IPv4)—You may not have enough IPv4 addresses to accommodate the number of IPv6 addresses. We recommend using a dynamic PAT pool to provide a large number of IPv4 translations.

Additional guidelines for NAT

- NAT rules apply to through the device traffic only, they do not apply to traffic initiated by the device, such as a RADIUS authentication.
- For interfaces that are members of a bridge group, you write NAT rules for the member interfaces. You cannot write NAT rules for the Bridge Virtual Interface (BVI) itself.
- You cannot write NAT rules for a Virtual Tunnel Interface (VTI), which are used in site-to-site VPN. Writing rules for the VTI's source interface will not apply NAT to the VPN tunnel. To write NAT rules that will apply to VPN traffic tunneled on a VTI, you must use "any" as the interface; you cannot explicitly specify interface names.

- (Network object NAT only.) You can only define a single NAT rule for a given object; if you want to configure multiple NAT rules for an object, you need to create multiple objects with different names that specify the same IP address.
- If a VPN is defined on an interface, inbound ESP traffic on the interface is not subject to the NAT rules. The system allows the ESP traffic for established VPN tunnels only, dropping traffic not associated with an existing tunnel. This restriction applies to ESP and UDP ports 500 and 4500.
- If you define a site-to-site VPN on a device that is behind a device that is applying dynamic PAT, so that UDP ports 500 and 4500 are not the ones actually used, you must initiate the connection from the device that is behind the PAT device. The responder cannot initiate the security association (SA) because it does not know the correct port numbers.
- If you change the NAT configuration, and you do not want to wait for existing translations to time out before the new NAT configuration is used, you can clear the translation table using the **clear xlate** command in the device CLI. However, clearing the translation table disconnects all current connections that use translations.

  If you create a new NAT rule that should apply to an existing connection (such as a VPN tunnel), you need to use **clear conn** to end the connection. Then, the attempt to re-establish the connection should hit the NAT rule and the connection should be NAT'ed correctly.



Note If you remove a dynamic NAT or PAT rule, and then add a new rule with mapped addresses that overlap the addresses in the removed rule, then the new rule will not be used until all connections associated with the removed rule time out or are cleared using the **clear xlate** or **clear conn** commands. This safeguard ensures that the same address is not assigned to multiple hosts.

- When translating SCTP traffic, use static network object NAT only. Dynamic NAT/PAT is not allowed. Although you can configure static twice NAT, this is not recommended because the topology of the destination part of the SCTP association is unknown.
- Objects and object groups used in NAT cannot be undefined; they must include IP addresses.
- You cannot use an object group with both IPv4 and IPv6 addresses; the object group must include only one type of address.
- (Twice NAT only.) When using **any** as the source address in a NAT rule, the definition of “any” traffic (IPv4 vs. IPv6) depends on the rule. Before the ASA performs NAT on a packet, the packet must be IPv6-to-IPv6 or IPv4-to-IPv4; with this prerequisite, the ASA can determine the value of **any** in a NAT rule. For example, if you configure a rule from “any” to an IPv6 server, and that server was mapped from an IPv4 address, then **any** means “any IPv6 traffic.” If you configure a rule from “any” to “any,” and you map the source to the interface IPv4 address, then **any** means “any IPv4 traffic” because the mapped interface address implies that the destination is also IPv4.
- You can use the same mapped object or group in multiple NAT rules.
- The mapped IP address pool cannot include:
  - The mapped interface IP address. If you specify “any” interface for the rule, then all interface IP addresses are disallowed. For interface PAT (routed mode only), specify the interface name instead of the interface address.

  - The failover interface IP address.
  - (Transparent mode.) The management IP address.
  - (Dynamic NAT.) The standby interface IP address when VPN is enabled.
  - Existing VPN pool addresses.

- Avoid using overlapping addresses in static and dynamic NAT policies. For example, with overlapping addresses, a PPTP connection can fail to get established if the secondary connection for PPTP hits the static instead of dynamic xlate.
- You cannot use overlapping addresses in the source address of a NAT rule and a remote access VPN address pool.
- For application inspection limitations with NAT or PAT, see Default Inspections and NAT Limitations, on page 302.
- The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. See Routing NAT Packets, on page 243 for more information.
- If you enable the **arp permit-nonconnected** command, the system does not respond to ARP requests if the mapped address is not part of any connected subnet and you also do not specify the mapped interface in the NAT rule (that is, you specify "any" interface). To resolve this problem, specify the mapped interface.
- If you specify a destination interface in a rule, then that interface is used as the egress interface rather than looking up the route in the routing table. However, for identity NAT, you have the option to use a route lookup instead.
- If you use PAT on Sun RPC traffic, which is used to connect to NFS servers, be aware that the NFS server might reject connections if the PAT'ed port is above 1024. The default configuration of NFS servers is to reject connections from ports higher than 1024. The error is typically "Permission Denied." Mapping ports above 1024 might happen if you use the "flat range" option to use the higher port numbers if a port in the lower range is not available, especially if you do not select the option to include the lower range in the flat range.Mapping ports above 1024 happens if you do not select the option to include the reserved ports (1-1023) in the port range of a PAT pool. You can avoid this problem by changing the NFS server configuration to allow all port numbers.
- NAT applies to through traffic only. Traffic generated by the system is not subject to NAT.
- You can improve system performance and reliability by using the transactional commit model for NAT. See the basic settings chapter in the general operations configuration guide for more information. The option is under **Configurations > Device Management > Advanced > Rule Engine**.
- Do not name a network object or group pat-pool, using any combination of upper- or lower-case letters.
- The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.
- You cannot use NAT on the internal payload of Protocol Independent Multicast (PIM) registers.
- (Twice NAT) When writing NAT rules for a dual ISP interface setup (primary and backup interfaces using service level agreements in the routing configuration), do not specify destination criteria in the rule. Ensure the rule for the primary interface comes before the rule for the backup interface. This allows the device to choose the correct NAT destination interface based on the current routing state when the

primary ISP is unavailable. If you specify destination objects, the NAT rule will always select the primary interface for the otherwise duplicate rules.

- If you get the ASP drop reason nat-no-xlate-to-pat-pool for traffic that should not match the NAT rules defined for the interface, configure identity NAT rules for the affected traffic so the traffic can pass untranslated.
- If you configure NAT for GRE tunnel endpoints, you must disable keepalives on the endpoints or the tunnel cannot be established. The endpoints send keepalives to the original addresses.
- DHCP and BOOTP share ports UDP/67-68. Because BOOTP is obsolete, writing NAT rules for the bootps port can cause port allocation problems when also running DHCP. Consider using DHCP relay instead for transmitting DHCP requests between network segments.
- In rare cases, return traffic (server to client) with an existing translation (xlate) might be logged as a new flow in Connection Events. This can occur when the client has already terminated the connection and the server sends another packet that reaches the device during the brief interval between connection closure and xlate removal, often due to application behavior or TCP stack cleanup. Because the device removes the xlate only after deleting the connection, a server packet can arrive while the xlate still exists. If no valid connection entry is found, the device logs a separate connection event based on the matched Access Control Policy rule.

Network Object NAT Guidelines for Mapped Address Objects

For dynamic NAT, you must use an object or group for the mapped addresses. For the other NAT types, you can use an object or group, or you have the option of using inline addresses. Network object groups are particularly useful for creating a mapped address pool with discontinuous IP address ranges or multiple hosts or subnets.

Consider the following guidelines when creating objects for mapped addresses.

- A network object group can contain objects or inline addresses of either IPv4 or IPv6 addresses. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.
- See Additional guidelines for NAT, on page 156 for information about disallowed mapped IP addresses.
- Do not name a network object or group pat-pool, using any combination of upper- or lower-case letters.
- Dynamic NAT:
  - You cannot use an inline address; you must configure a network object or group.
  - The object or group cannot contain a subnet; the object must define a range; the group can include hosts and ranges.
  - If a mapped network object contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.
- Dynamic PAT (Hide):
  - Instead of using an object, you can optionally configure an inline host address or specify the interface address.
  - If you use an object, the object or group cannot contain a subnet. The object must define a host, or for a PAT pool, a range. The group (for a PAT pool) can include hosts and ranges.

- Static NAT or Static NAT with port translation:
  - Instead of using an object, you can configure an inline address or specify the interface address (for static NAT-with-port-translation).
  - If you use an object, the object or group can contain a host, range, or subnet.
- Identity NAT
  - Instead of using an object, you can configure an inline address.
  - If you use an object, the object must match the real addresses you want to translate.

Twice NAT Guidelines for Real and Mapped Address Objects

For each NAT rule, configure up to four network objects or groups for:

- Source real address
- Source mapped address
- Destination real address
- Destination mapped address

Objects are required unless you specify the **any** keyword inline to represent all traffic, or for some types of NAT, the **interface** keyword to represent the interface address. Network object groups are particularly useful for creating a mapped address pool with discontinuous IP address ranges or multiple hosts or subnets.

Consider the following guidelines when creating objects for twice NAT.

- A network object group can contain objects or inline addresses of either IPv4 or IPv6 addresses. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.
- See Additional guidelines for NAT, on page 156 for information about disallowed mapped IP addresses.
- Do not name a network object or group pat-pool, using any combination of upper- or lower-case letters.
- Source Dynamic NAT:
  - You typically configure a larger group of real addresses to be mapped to a smaller group.
  - The mapped object or group cannot contain a subnet; the object must define a range; the group can include hosts and ranges.
  - If a mapped network object contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.
- Source Dynamic PAT (Hide):
  - If you use an object, the object or group cannot contain a subnet. The object must define a host, or for a PAT pool, a range. The group (for a PAT pool) can include hosts and ranges.
- Source Static NAT or Static NAT with port translation:
  - The mapped object or group can contain a host, range, or subnet.

    - The static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired.

- Source Identity NAT
    - The real and mapped objects must match. You can use the same object for both, or you can create separate objects that contain the same IP addresses.

- Destination Static NAT or Static NAT with port translation (the destination translation is always static):
    - Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 152.
    - For identity NAT, the real and mapped objects must match. You can use the same object for both, or you can create separate objects that contain the same IP addresses.
    - The static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired.
    - For static interface NAT with port translation (routed mode only), you can specify the **interface** keyword instead of a network object/group for the mapped address.

Twice NAT Guidelines for Service Objects for Real and Mapped Ports

You can optionally configure service objects for:

- **Source real port (Static only)** *or* **Destination real port**
- **Source mapped port (Static only)** *or* **Destination mapped port**

Consider the following guidelines when creating objects for twice NAT.

- NAT supports TCP, UDP, and SCTP only. When translating a port, be sure the protocols in the real and mapped service objects are identical (for example, both TCP). Although you can configure static twice NAT rules with SCTP port specifications, this is not recommended, because the topology of the destination part of the SCTP association is unknown. Use static object NAT instead for SCTP.
- The "not equal" (**neq**) operator is not supported.
- For identity port translation, you can use the same service object for both the real and mapped ports.
- Source Dynamic NAT—Source Dynamic NAT does not support port translation.
- Source Dynamic PAT (Hide)—Source Dynamic PAT does not support port translation.
- Source Static NAT, Static NAT with port translation, or Identity NAT—A service object can contain both a source and destination port; however, you should specify *either* the source *or* the destination port for both service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. For example, if you want to translate the port for the source host, then configure the source service.

- Destination Static NAT or Static NAT with port translation (the destination translation is always static)—For non-static source NAT, you can only perform port translation on the destination. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored.

Dynamic NAT

The following topics explain dynamic NAT and how to configure it.

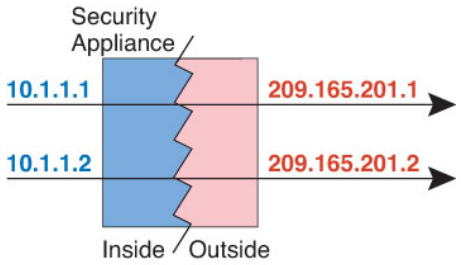
About dynamic NAT

Dynamic NAT translates a group of real addresses to a pool of mapped addresses that are routable on the destination network. The mapped pool typically includes fewer addresses than the real group. When a host you want to translate accesses the destination network, NAT assigns the host an IP address from the mapped pool. The translation is created only when the real host initiates the connection. The translation is in place only for the duration of the connection, and a given user does not keep the same IP address after the translation times out. Users on the destination network, therefore, cannot initiate a reliable connection to a host that uses dynamic NAT, even if the connection is allowed by an access rule.

Note For the duration of the translation, a remote host can initiate a connection to the translated host if an access rule allows it. Because the address is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule. A successful connection from a remote host can reset the idle timer for the connection.

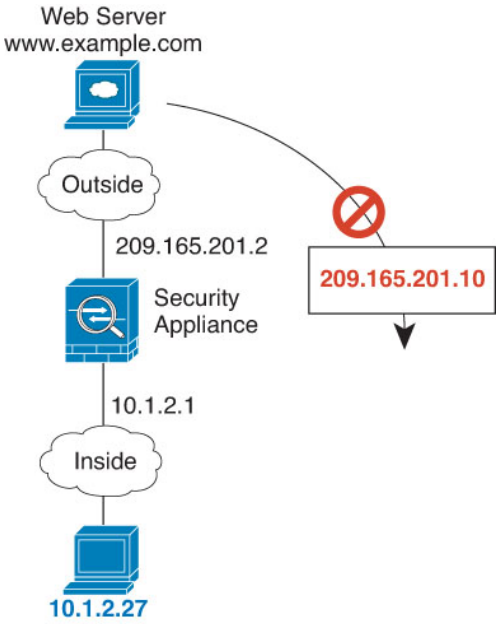
The following figure shows a typical dynamic NAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back.

Figure 18: Dynamic NAT



The following figure shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table; therefore, the packet is dropped.

Figure 19: Remote Host Attempts to Initiate a Connection to a Mapped Address



Dynamic NAT disadvantages and advantages

Dynamic NAT has these disadvantages:

- If the mapped pool has fewer addresses than the real group, you could run out of addresses if the amount of traffic is more than expected.

  Use PAT or a PAT fall-back method if this event occurs often because PAT provides over 64,000 translations using ports of a single address.

- You have to use a large number of routable addresses in the mapped pool, and routable addresses may not be available in large quantities.

The advantage of dynamic NAT is that some protocols cannot use PAT. PAT does not work with the following:

- IP protocols that do not have a port to overload, such as GRE version 0.
- Some multimedia applications that have a data stream on one port, the control path on another port, and are not open standard.

Configure Dynamic Network Object NAT

This section describes how to configure network object NAT for dynamic NAT.

Procedure

Step 1 Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration** > **Firewall** > **NAT Rules**, then click **Add > Add Network Object NAT Rule**.
- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then edit a network object.

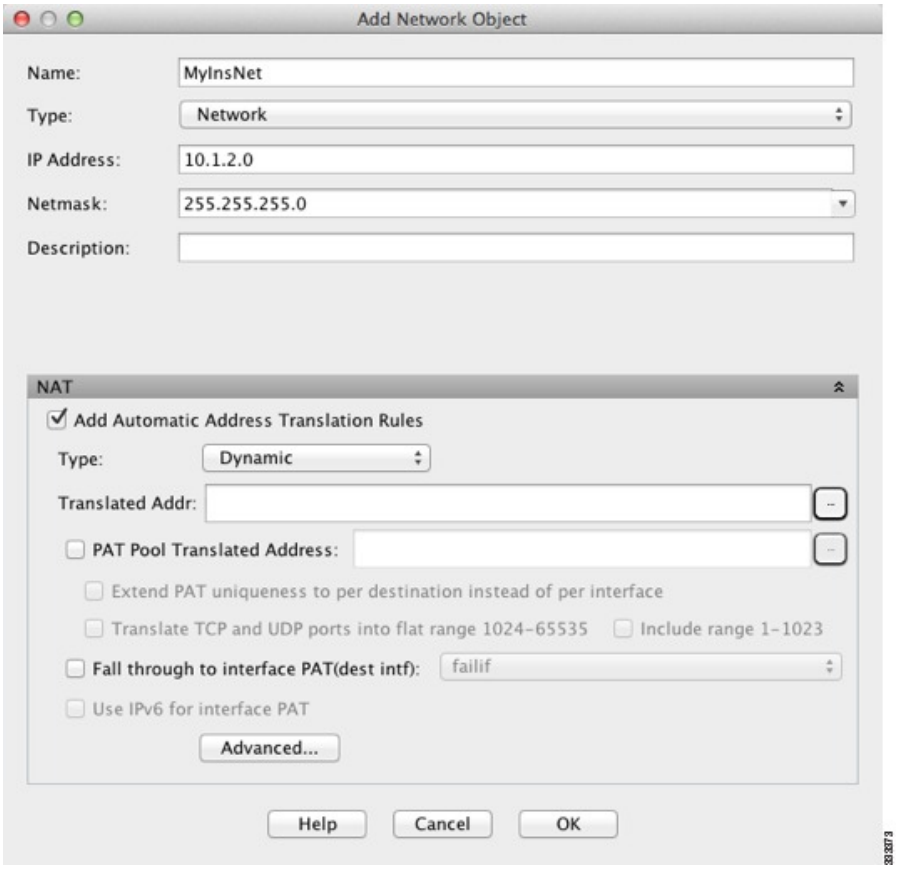
Step 2 For a new object, enter values for the following fields:

a) **Name**—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.

b) **Type**—Host, Network, or Range.

c) **IP Addresses**—IPv4 or IPv6 addresses, a single address for a host, a starting and ending address for a range, and for subnet, either an IPv4 network address and mask (for example, 10.100.10.0 255.255.255.0) or IPv6 address and prefix length (for example, 2001:DB8:0:CD30::/60).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.



Step 4 Check the **Add Automatic Translation Rules** check box.

Step 5 From the Type drop-down list, choose **Dynamic**.

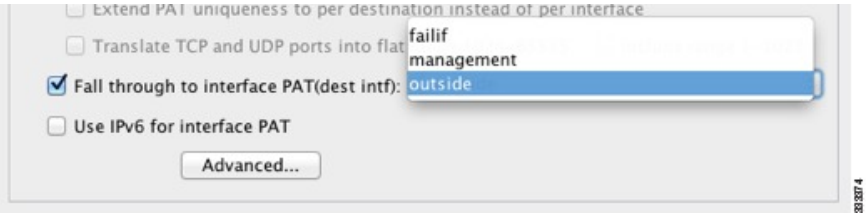
Step 6 To the right of the Translated Addr field, click the browse button and choose the network object or network object group that contains the mapped addresses.

You can create a new object if necessary.

The object or group cannot contain a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

If a mapped network object contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.

Step 7 (Optional, when mapped interface is a non-bridge group member only.) To use the interface IP address as a backup method when the other mapped addresses are already allocated, check the **Fall through to interface PAT (dest intf)** check box, and choose the interface from the drop-down list. To use the IPv6 address of the interface, also check the **Use IPv6 for interface PAT** check box.



Step 8 (Optional) Click **Advanced**, configure the following options in the Advanced NAT Settings dialog box, and click **OK**.

- **Translate DNS replies for rule**—Translates the IP address in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). See Rewriting DNS queries and responses using NAT, on page 262 for more information.
- (Required for bridge group member interfaces.) **Interface**—Specifies the real interface (**Source**) and the mapped interface (**Destination**) where this NAT rule applies. By default, the rule applies to all interfaces except for bridge group members.

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

Configure Dynamic Twice NAT

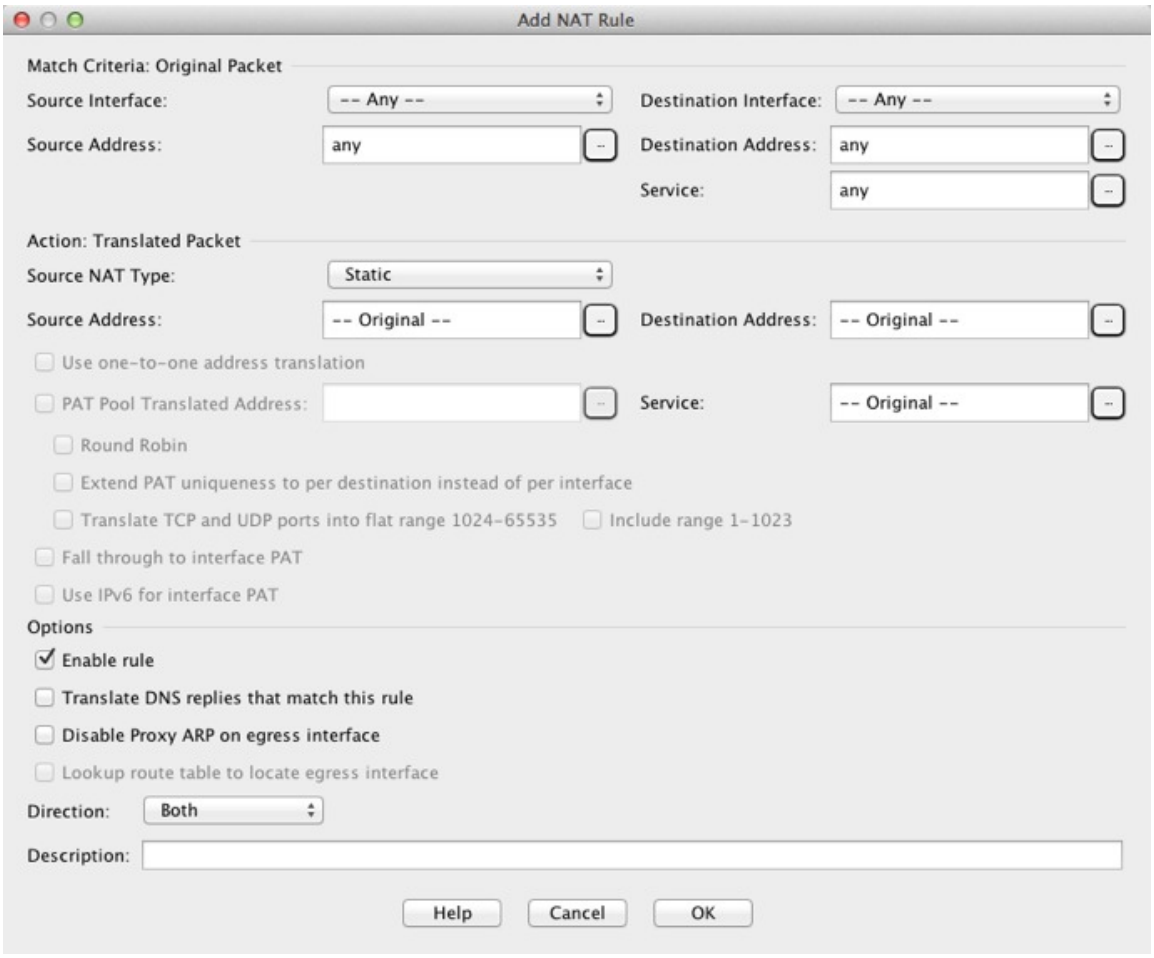
This section describes how to configure twice NAT for dynamic NAT.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**, and then do one of the following:

- Click **Add**, or **Add > Add NAT Rule Before Network Object NAT Rules**.
- Click **Add > Add NAT Rule After Network Object NAT Rules**.
- Select a twice NAT rule and click **Edit**.

The Add NAT Rule dialog box appears.

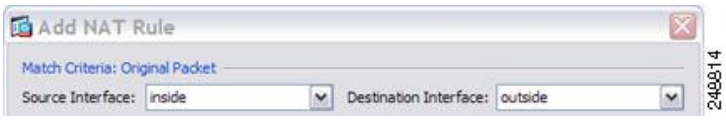


Step 2 (Required for bridge group member interfaces.) Set the source and destination interfaces.

In routed mode, the default is any interface for both source and destination. You can select specific interfaces for one or both options. However, you must select the interface when writing a rule for bridge group member interfaces; "any" does not include these interfaces.

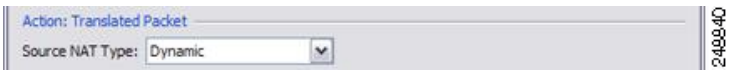
a) From the **Match Criteria: Original Packet > Source Interface** drop-down list, choose the source interface.

b) From the **Match Criteria: Original Packet > Destination Interface** drop-down list, choose the destination interface.

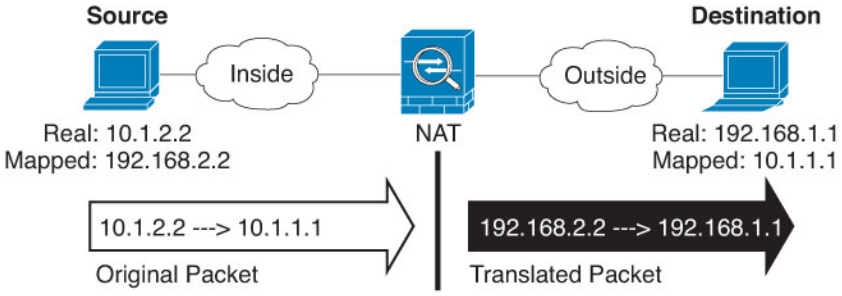


Step 3 Choose **Dynamic** from the **Action: Translated Packet** > **Source NAT Type** drop-down list.

This setting only applies to the source address; the destination translation is always static.



Step 4 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.



a) For the **Match Criteria: Original Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**.

b) (Optional.) For the **Match Criteria: Original Packet > Destination Address**, click the browse button and choose an existing network object, group, or interface, or create a new object or group from the Browse Original Destination Address dialog box. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 152.

For static interface NAT with port translation only, choose an interface from the Browse dialog box. Be sure to also configure a service translation. For this option, you must configure a specific interface for the Source Interface. See Static NAT with port translation, on page 190 for more information.

Step 5 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired.

a) For **Action: Translated Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated Source Address dialog box.

For dynamic NAT, you typically configure a larger group of source addresses to be mapped to a smaller group.

If a mapped network object contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.

Note
The object or group cannot contain a subnet.

b) For **Action: Translated Packet > Destination Address**, click the browse button and choose an existing network object or group, or create a new object or group from the Browse Translated Destination Address dialog box.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

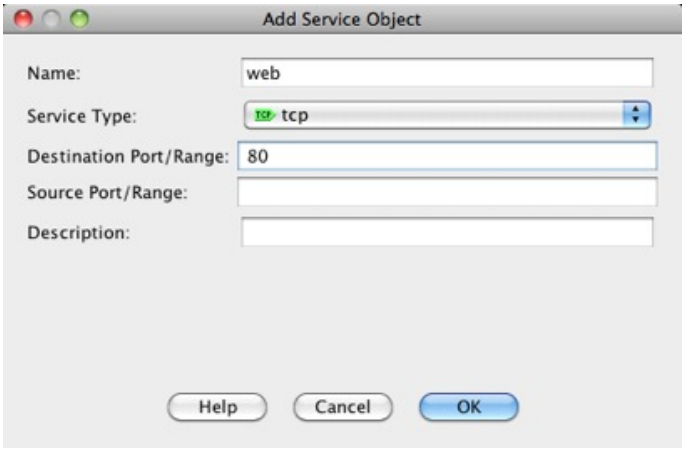
If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see Static NAT, on page 189. See Additional guidelines for NAT, on page 156 for information about disallowed mapped IP addresses.

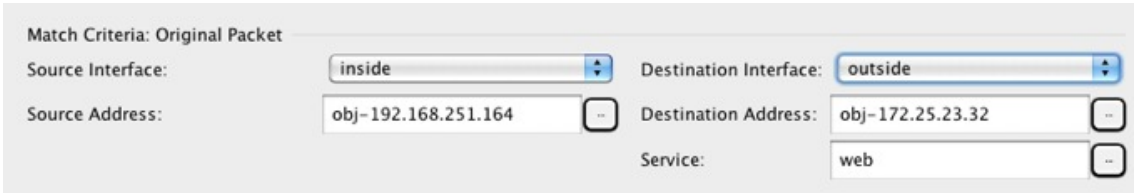
Step 6 (Optional.) Identify the destination service ports for service translation.

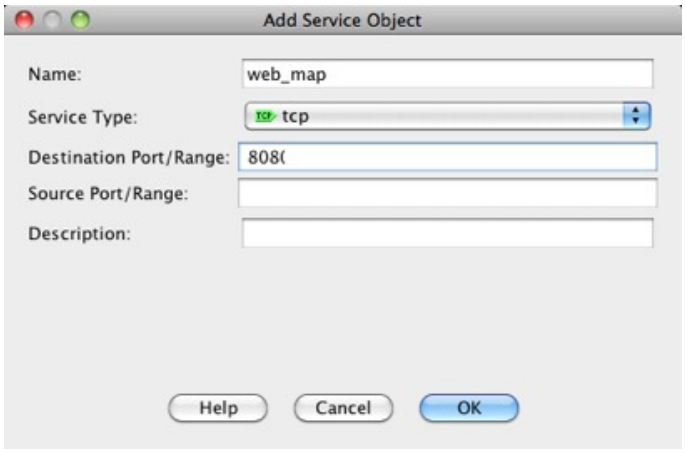
- Identify the original packet port (the *mapped destination port*). For **Match Criteria: Original Packet > Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Original Service dialog box.
- Identify the translated packet port (the *real destination port*). For **Action: Translated Packet > Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Translated Service dialog box.

Dynamic NAT does not support port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

For example:



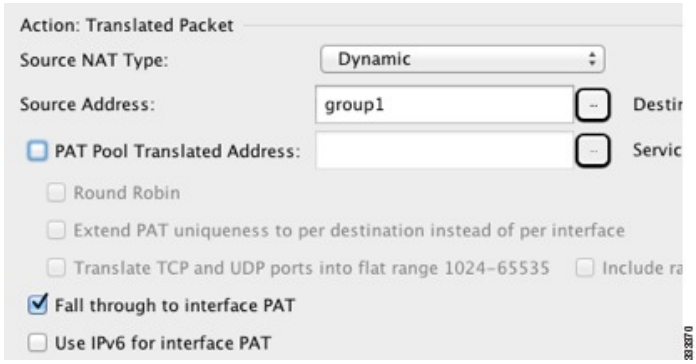




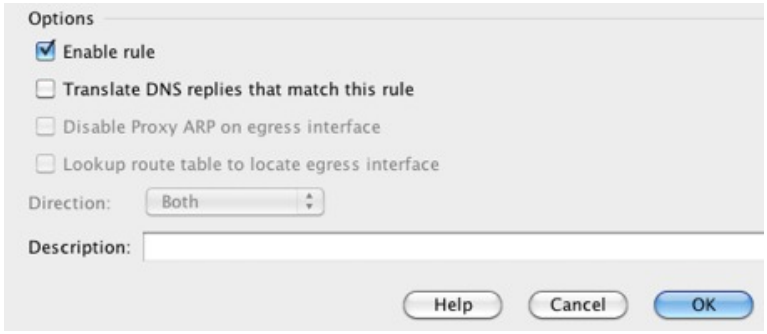


Step 7 (Optional, when mapped interface is a non-bridge group member only.) To use the interface IP address as a backup method if the other mapped source addresses are already allocated, check the **Fall through to interface PAT** check box. To use the IPv6 interface address, also check the **Use IPv6 for interface PAT** check box.

The destination interface IP address is used. This option is only available if you configure a specific Destination Interface.



Step 8 (Optional.) Configure NAT options in the Options area.



- **Enable rule** —Enables this NAT rule. The rule is enabled by default.
- (For a source-only rule) **Translate DNS replies that match this rule**—Rewrites the DNS A record in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). You cannot configure DNS modification if you configure a destination address. See Rewriting DNS queries and responses using NAT, on page 262 for more information.
- **Description**—Adds a description about the rule up to 200 characters in length.

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

Dynamic PAT

The following topics describe dynamic PAT.

About dynamic PAT

Dynamic PAT translates multiple real addresses to a single mapped IP address by translating the real address and source port to the mapped address and a unique port.

Each connection requires a separate translation session because the source port differs for each connection. For example, 10.1.1.1:1025 requires a separate translation from 10.1.1.1:1026.

The following figure shows a typical dynamic PAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back. The mapped address is the same for each translation, but the port is dynamically assigned.

Figure 20: Dynamic PAT



For the duration of the translation, a remote host on the destination network can initiate a connection to the translated host if an access rule allows it. Because the port address (both real and mapped) is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.

After the connection expires, the port translation also expires. For multi-session PAT, the PAT timeout is used, 30 seconds by default. For per-session PAT, the xlate is immediately removed.

Note We recommend that you use different PAT pools for each interface. If you use the same pool for multiple interfaces, especially if you use it for "any" interface, the pool can be quickly exhausted, with no ports available for new translations.

Dynamic PAT disadvantages and advantages

Dynamic PAT lets you use a single mapped address, thus conserving routable addresses. You can even use the ASA interface IP address as the PAT address.

You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.

Dynamic PAT does not work with some multimedia applications that have a data stream that is different from the control path. See Default Inspections and NAT Limitations, on page 302 for more information about NAT and PAT support.

Dynamic PAT might also create a large number of connections appearing to come from a single IP address, and servers might interpret the traffic as a DoS attack. You can configure a PAT pool of addresses and use a round-robin assignment of PAT addresses to mitigate this situation.

PAT pool object guidelines

When creating network objects for a PAT pool, follow these guidelines.

For a PAT pool

- Ports are mapped to an available port in the 1024 to 65535 range. You can optionally include the reserved ports, those below 1024, to make the entire port range available for translations.

  When operating in a cluster, blocks of 512 ports per address are allocated to the members of the cluster, and mappings are made within these port blocks. If you also enable block allocation, the ports are distributed according to the block allocation size, whose default is also 512. If you change the cluster unit limit (the size of the cluster), ensure that you clear xlates or reboot the devices so that PAT pools can be appropriately reallocated to the cluster units.

- If you enable block allocation for a PAT pool, port blocks are allocated in the 1024-65535 range only. Thus, if an application requires a low port number (1-1023), it might not work. For example, an application requesting port 22 (SSH) will get a mapped port within the range of 1024-65535 and within the block allocated to the host.

- If you use the same PAT pool object in two separate rules, then be sure to specify the same options for each rule. For example, if one rule specifies extended PAT, then the other rule must also specify extended PAT.

- If a host has an existing connection, then subsequent connections from that host use the same PAT IP address. If no ports are available, this can prevent the connection. Use the round robin option to avoid this problem.

- For best performance, limit the number of IP addresses within a PAT pool to 10,000.

For extended PAT for a PAT pool

- Many application inspections do not support extended PAT. See Default Inspections and NAT Limitations, on page 302 for a complete list of unsupported inspections.

- If you enable extended PAT for a dynamic PAT rule, then you cannot also use an address in the PAT pool as the PAT address in a separate static NAT with port translation rule. For example, if the PAT pool includes 10.1.1.1, then you cannot create a static NAT-with-port-translation rule using 10.1.1.1 as the PAT address.

- If you use a PAT pool and specify an interface for fallback, you cannot specify extended PAT.
- For VoIP deployments that use ICE or TURN, do not use extended PAT. ICE and TURN rely on the PAT binding to be the same for all destinations.
- You cannot use extended PAT on units in a cluster.
- Extended PAT increases memory usage on the device.

For round robin for a PAT pool

- If a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. However, this "stickiness" does not survive a failover. If the device fails over, then subsequent connections from a host might not use the initial IP address.
- IP address "stickiness" is also impacted if you mix PAT pool/round robin rules with interface PAT rules on the same interface. For any given interface, choose either a PAT pool or interface PAT; do not create competing PAT rules.
- Round robin, especially when combined with extended PAT, can consume a large amount of memory. Because NAT pools are created for every mapped protocol/IP address/port range, round robin results in a large number of concurrent NAT pools, which use memory. Extended PAT results in an even larger number of concurrent NAT pools.

Configure Dynamic Network Object PAT (Hide)

This section describes how to configure network object NAT for dynamic PAT (hide), which uses a single address for translation instead of a PAT pool.

Procedure

Step 1 Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration** > **Firewall** > **NAT Rules**, then click **Add > Add Network Object NAT Rule**.
- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then edit a network object.

Step 2 For a new object, enter values for the following fields:

a) **Name**—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.

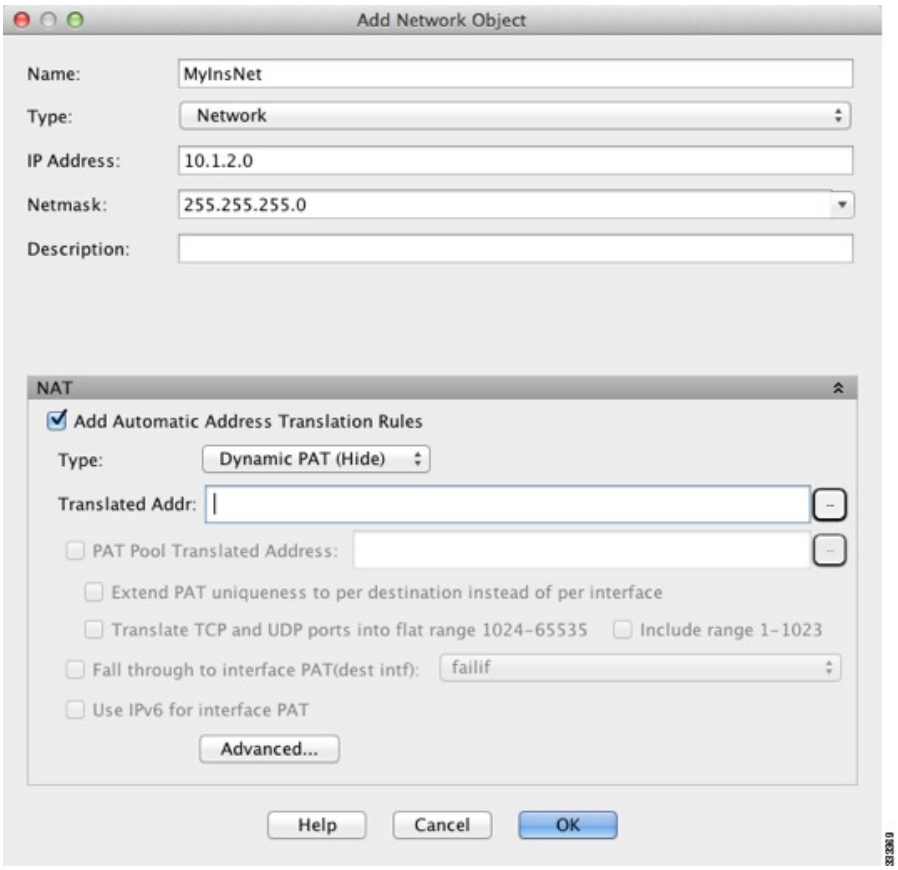
b) **Type**—Host, Network, or Range.

c) **IP Addresses**—IPv4 or IPv6 addresses, a single address for a host, a starting and ending address for a range, and for subnet, either an IPv4 network address and mask (for example, 10.100.10.0 255.255.255.0) or IPv6 address and prefix length (for example, 2001:DB8:0:CD30::/60).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

Step 4 Check the **Add Automatic Translation Rules** check box.

Step 5 From the Type drop-down list, choose **Dynamic PAT (Hide)**.



Step 6 Specify a single mapped address. In the Translated Addr. field, specify the mapped IP address by doing one of the following:

- Type a host IP address.
- Click the browse button and select a host network object (or create a new one).
- (Non-bridge group member interfaces only.) Type an interface name or click the browse button, and choose an interface from the Browse Translated Addr dialog box.



If you specify an interface name, then you enable *interface PAT*, where the specified interface IP address is used as the mapped address. To use the IPv6 interface address, you must also check the **Use IPv6 for interface PAT** check box. With interface PAT, the NAT rule only applies to the specified mapped interface, which cannot be a member of a bridge group. (If you do not use interface PAT, then the rule applies to all interfaces by default.) You cannot specify an interface in transparent mode.

Step 7 (Optional.) Click **Advanced**, configure the following options in the Advanced NAT Settings dialog box, and click **OK**.

- (Required for bridge group member interfaces.) **Interface**—Specifies the real interface (**Source**) and the mapped interface (**Destination**) where this NAT rule applies. By default, the rule applies to all interfaces except for bridge group members.

Step 8 Click **OK**, and then **Apply**.

Configure Dynamic Network Object PAT Using a PAT Pool

This section describes how to configure network object NAT for dynamic PAT using a PAT pool.

Procedure

Step 1 Add NAT to a new or existing network object:

- To add a new network object NAT rule, choose **Configuration** > **Firewall** > **NAT Rules**, then click **Add** > **Add Network Object NAT Rule**.
- To add NAT to an existing network object, choose **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups**, and then edit a network object.

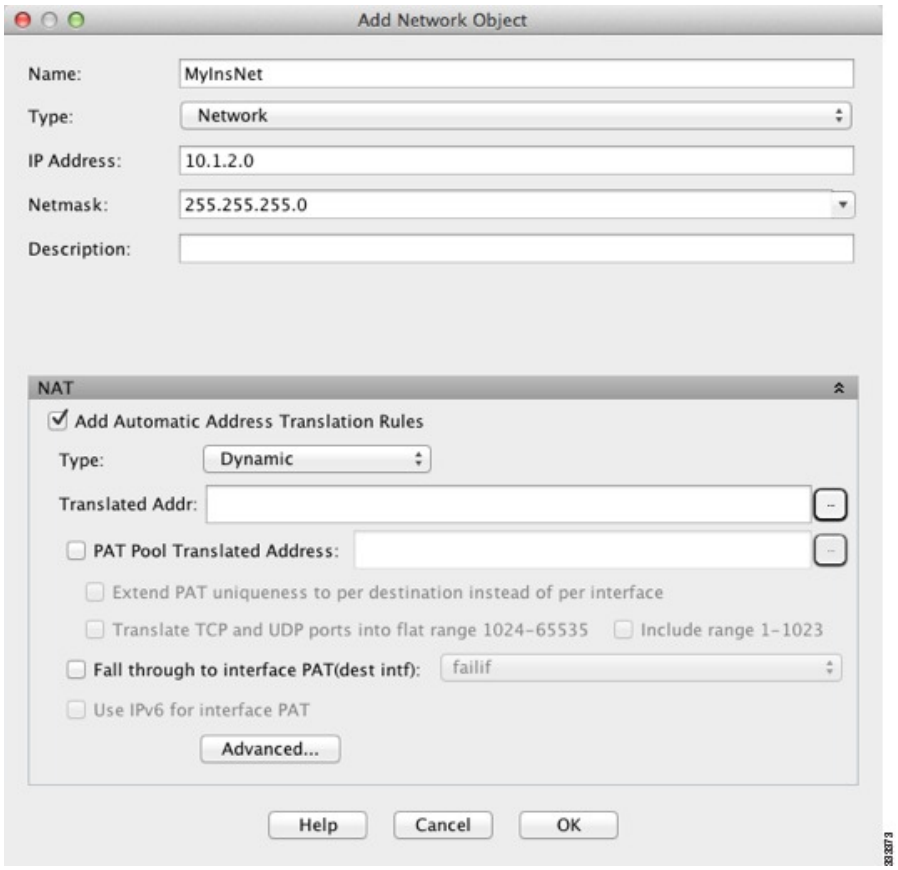
Step 2 For a new object, enter values for the following fields:

a) **Name**—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.

b) **Type**—Host, Network, or Range.

c) **IP Addresses**—IPv4 or IPv6 addresses, a single address for a host, a starting and ending address for a range, and for subnet, either an IPv4 network address and mask (for example, 10.100.10.0 255.255.255.0) or IPv6 address and prefix length (for example, 2001:DB8:0:CD30::/60).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.



Step 4 Check the **Add Automatic Translation Rules** check box.

Step 5 From the Type drop-down list, choose **Dynamic** even though you are configuring dynamic PAT with a PAT pool.

Step 6 To configure the PAT pool:

a) Do not enter a value for the Translated Addr. field; leave it blank.

b) Check the **PAT Pool Translated Address** check box, then click the browse button and choose the network object or group that contains the PAT pool addresses. Or create a new object from the Browse Translated PAT Pool Address dialog box.

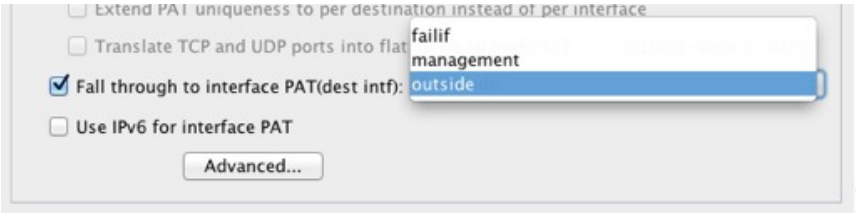
Note
The PAT pool object or group cannot contain a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.



c) (Optional) Select the following options as needed:

- **Round Robin**—To assign addresses and ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.
- **Extend PAT uniqueness to per destination instead of per interface** (8.4(3) and later, not including 8.5(1) or 8.6(1))—To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80.
- **Include Reserved Ports (1 to 1023)**—Includes the reserved ports, 1-1023, in the range of ports that are available for address translation. If you do not specify this option, addresses are translated to ports in the 1024-65535 range only.
- **Enable Block Allocation** (9.5.1 and later)—Enables port block allocation. For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time. If you allocate a block of ports, subsequent connections from the host use new randomly selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Port blocks are allocated in the 1024-65535 range only. Port block allocation is compatible with round robin, but you cannot use it with extended PAT. You also cannot use interface PAT fallback.

Step 7 (Optional, when mapped interface is a non-bridge group member only.) To use the interface IP address as a backup method when the other mapped addresses are already allocated, check the **Fall through to interface PAT** check box, and choose the interface from the drop-down list. To use the IPv6 address of the interface, also check the **Use IPv6 for interface PAT** check box.



Step 8 (Optional) Click **Advanced**, configure the following options in the Advanced NAT Settings dialog box, and click **OK**.

- (Required for bridge group member interfaces.) **Interface**—Specifies the real interface (**Source**) and the mapped interface (**Destination**) where this NAT rule applies. By default, the rule applies to all interfaces except for bridge group members.

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

Configure Dynamic Twice PAT (Hide)

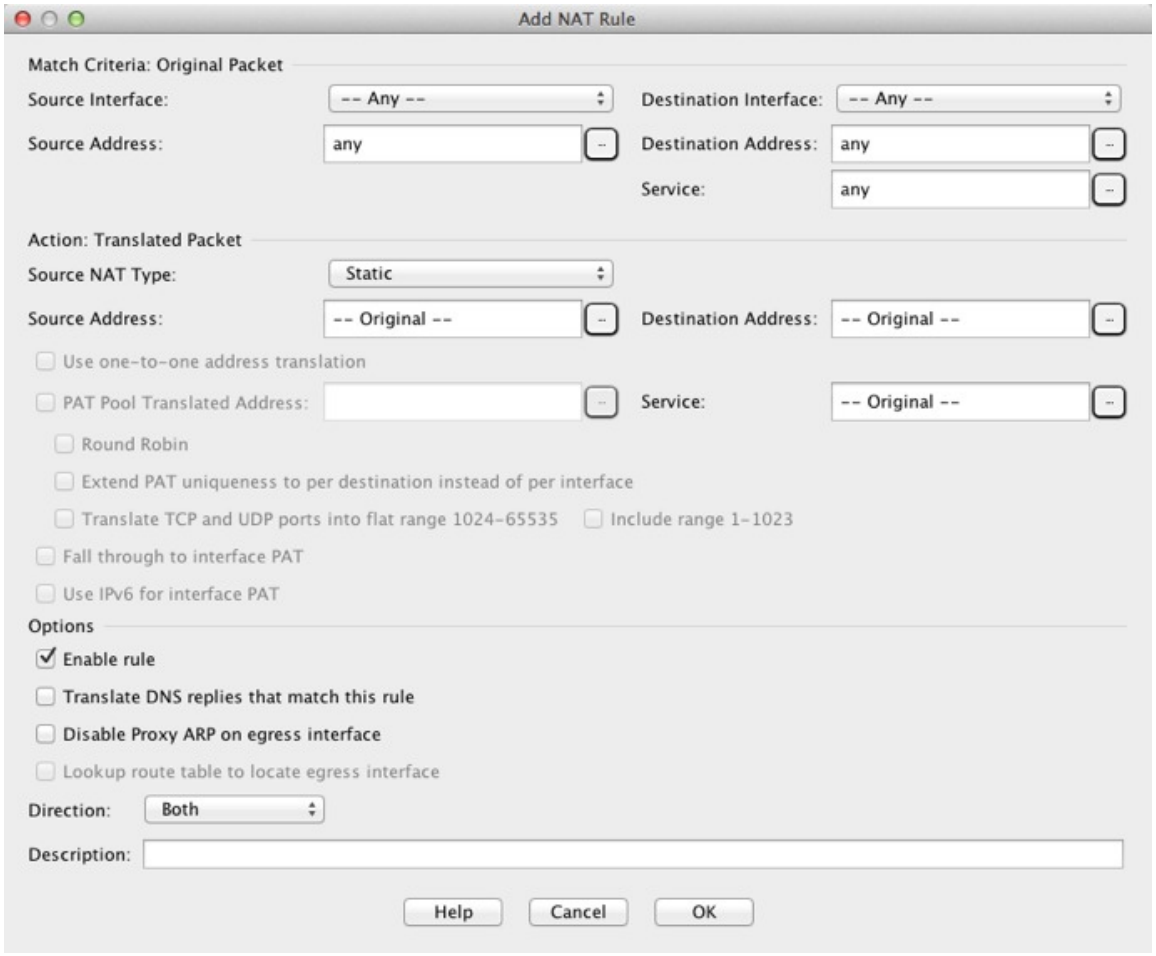
This section describes how to configure twice NAT for dynamic PAT (hide), which uses a single address for translation instead of a PAT pool.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**, and then do one of the following:

- Click **Add**, or **Add > Add NAT Rule Before Network Object NAT Rules**.
- Click **Add > Add NAT Rule After Network Object NAT Rules**.
- Select a twice NAT rule and click **Edit**.

The Add NAT Rule dialog box appears.

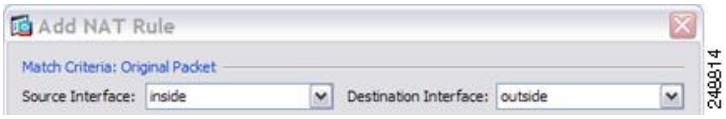


Step 2 (Required for bridge group member interfaces.) Set the source and destination interfaces.

In routed mode, the default is any interface for both source and destination. You can select specific interfaces for one or both options. However, you must select the interface when writing a rule for bridge group member interfaces; “any” does not include these interfaces.

a) From the **Match Criteria: Original Packet > Source Interface** drop-down list, choose the source interface.

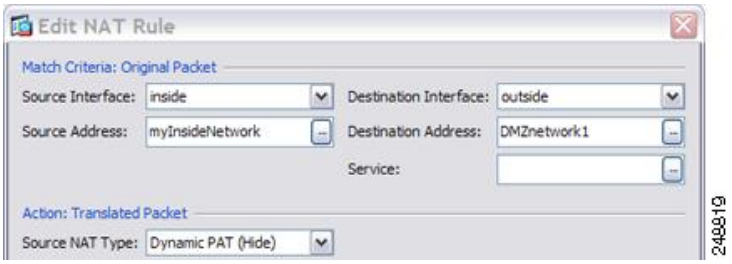
b) From the **Match Criteria: Original Packet > Destination Interface** drop-down list, choose the destination interface.



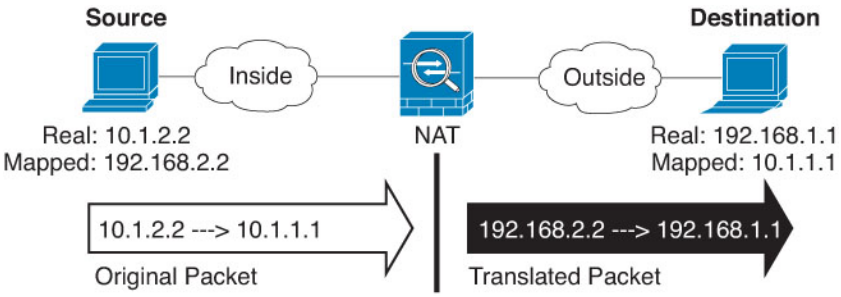
Step 3 Choose **Dynamic PAT (Hide)** from the **Action: Translated Packet** > **Source NAT Type** drop-down list.

This setting only applies to the source address; the destination translation is always static.

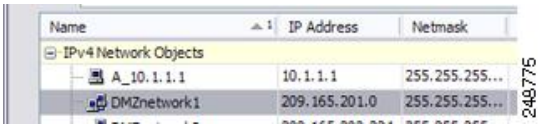
Note
To configure dynamic PAT using a PAT pool, choose **Dynamic** instead of Dynamic PAT (Hide), see Configure Dynamic Twice PAT Using a PAT Pool, on page 181.



Step 4 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.



a) For **Match Criteria: Original Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**.



b) (Optional) For **Match Criteria: Original Packet > Destination Address**, click the browse button and choose an existing network object, group, or interface (non-bridge group member interfaces only), or create a new object or group from the Browse Original Destination Address dialog box. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

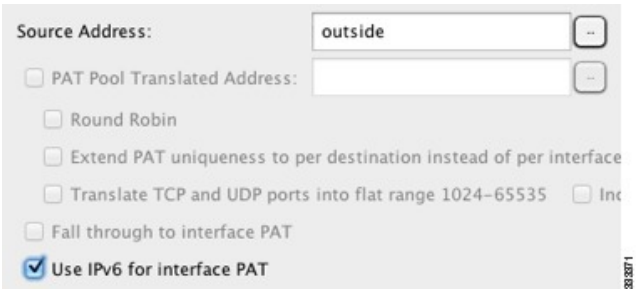
Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 152.

For static interface NAT with port translation only, choose an interface from the Browse dialog box. Be sure to also configure a service translation. For this option, you must configure a specific interface for the Source Interface. See Static NAT with port translation, on page 190 for more information.

Step 5 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired.

a) For **Action: Translated Packet > Source Address**, click the browse button and choose an existing network object that defines a host address, or an interface, or create a new object from the Browse Translated Source Address dialog box. The interface cannot be a bridge group member.

If you want to use the IPv6 address of the interface, check the **Use IPv6 for interface PAT** check box.



b) (Optional.) For **Action: Translated Packet > Destination Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated Destination Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

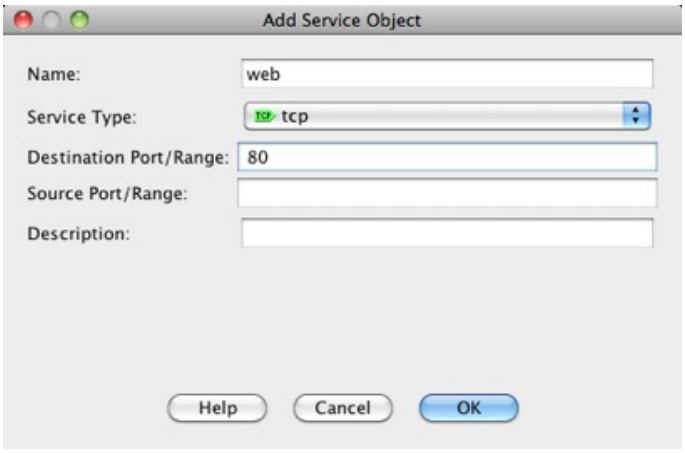
If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see Static NAT, on page 189. See Guidelines for NAT, on page 155 for information about disallowed mapped IP addresses.

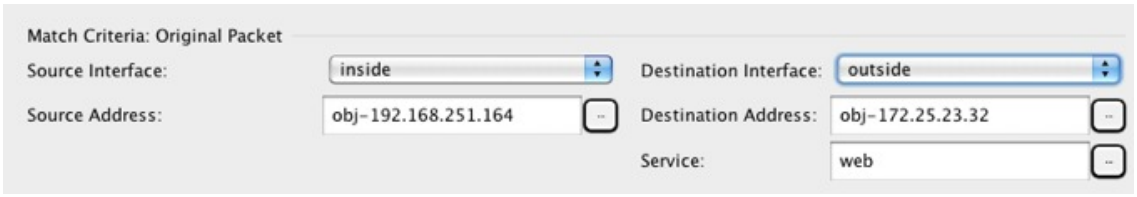
Step 6 (Optional.) Identify the destination service ports for service translation.

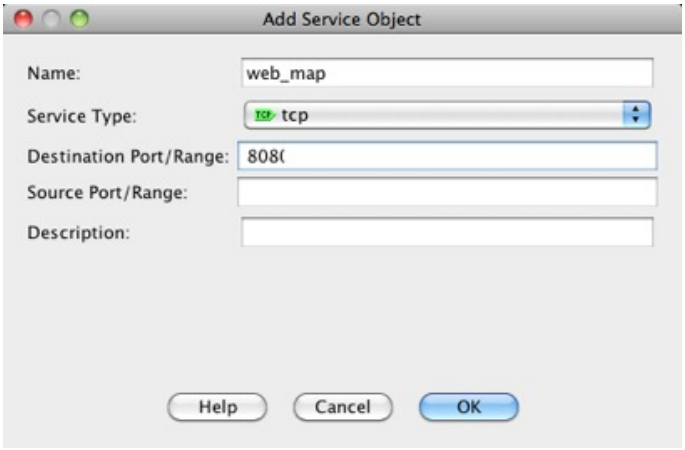
- Identify the original packet port (the *mapped destination port*). For **Match Criteria: Original Packet > Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Original Service dialog box.
- Identify the translated packet port (the *real destination port*). For **Action: Translated Packet > Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Translated Service dialog box.

Dynamic NAT does not support port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

For example:

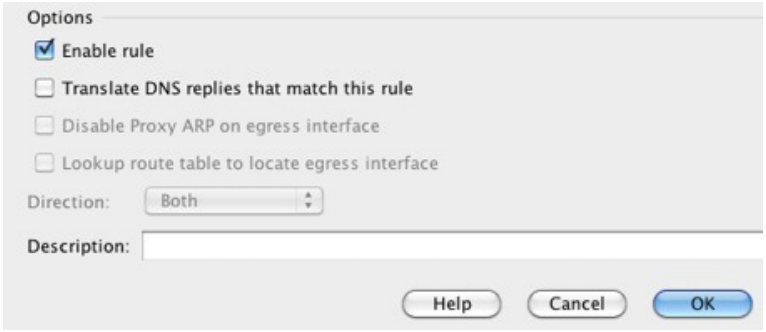








Step 7 (Optional) Configure NAT options in the Options area.



- **Enable rule**—Enables this NAT rule. The rule is enabled by default.
- **Description**—Adds a description about the rule up to 200 characters in length.

Step 8 Click **OK**, then click **Apply**.

Configure Dynamic Twice PAT Using a PAT Pool

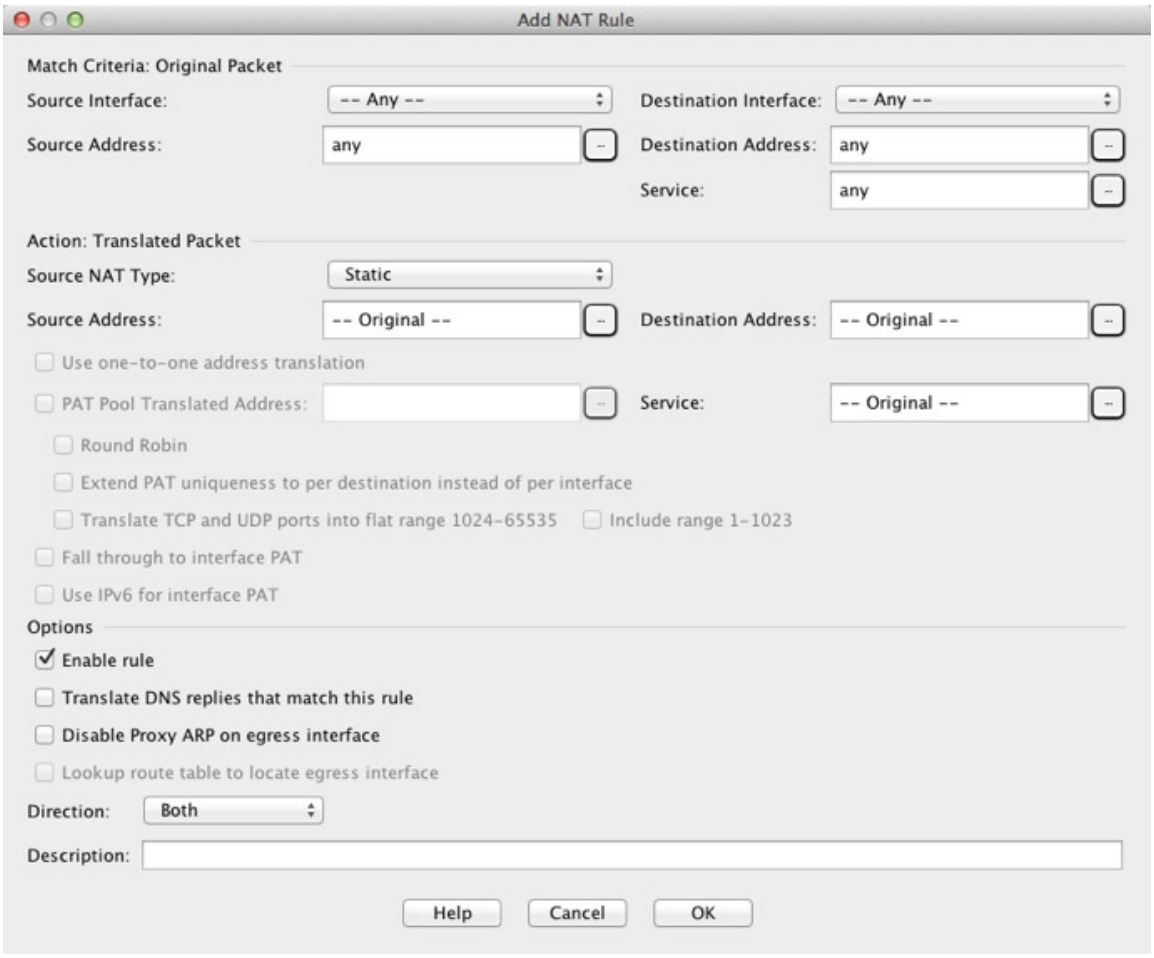
This section describes how to configure twice NAT for dynamic PAT using a PAT pool.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**, and then do one of the following:

- Click **Add**, or **Add** > **Add NAT Rule Before Network Object NAT Rules**.
- Click **Add** > **Add NAT Rule After Network Object NAT Rules**.
- Select a twice NAT rule and click **Edit**.

The Add NAT Rule dialog box appears.

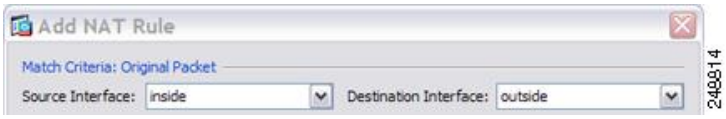


Step 2 (Required for bridge group member interfaces.) Set the source and destination interfaces.

In routed mode, the default is any interface for both source and destination. You can select specific interfaces for one or both options. However, you must select the interface when writing a rule for bridge group member interfaces; "any" does not include these interfaces.

a) From the **Match Criteria: Original Packet** > **Source Interface** drop-down list, choose the source interface.

b) From the **Match Criteria: Original Packet** > **Destination Interface** drop-down list, choose the destination interface.

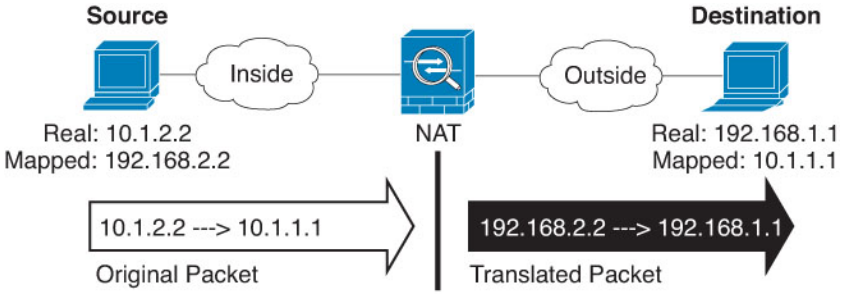


Step 3 Choose **Dynamic** from the **Action: Translated Packet** > **Source NAT Type** drop-down list.

This setting only applies to the source address; the destination translation is always static.



Step 4 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.



a) For the **Match Criteria: Original Packet** > **Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**.

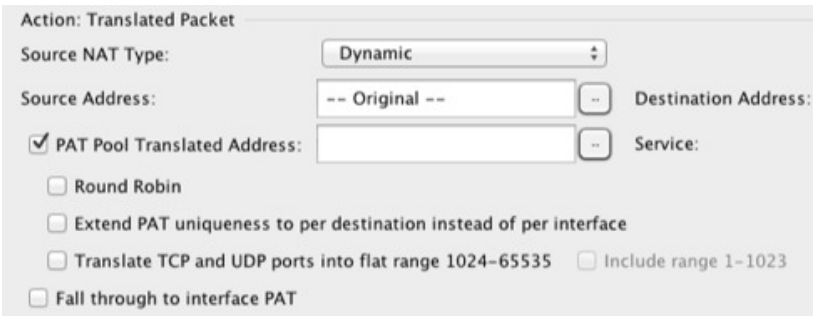
b) (Optional.) For the **Match Criteria: Original Packet** > **Destination Address**, click the browse button and choose an existing network object, group, or interface (non-bridge group member interfaces only), or create a new object or group from the Browse Original Destination Address dialog box. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 152.

For static interface NAT with port translation only, choose an interface from the Browse dialog box. Be sure to also configure a service translation. For this option, you must configure a specific interface for the Source Interface. See Static NAT with port translation, on page 190 for more information.

Step 5 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired.

a) Check the **PAT Pool Translated Address** check box, then click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated PAT Pool Address dialog box. **Note**: Leave the Source Address field empty.



Note

The object or group cannot contain a subnet.

b) (Optional.) For **Action: Translated Packet** > **Destination Address**, click the browse button and choose an existing network object or group, or create a new object or group from the Browse Translated Destination Address dialog box.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

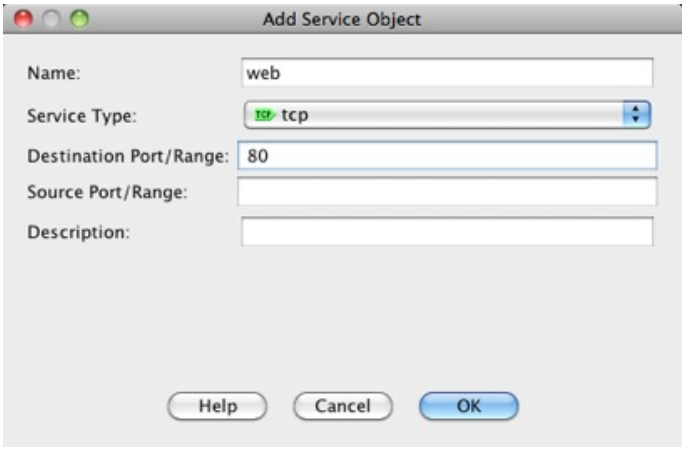
If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see Static NAT, on page 189. See Guidelines for NAT, on page 155 for information about disallowed mapped IP addresses.

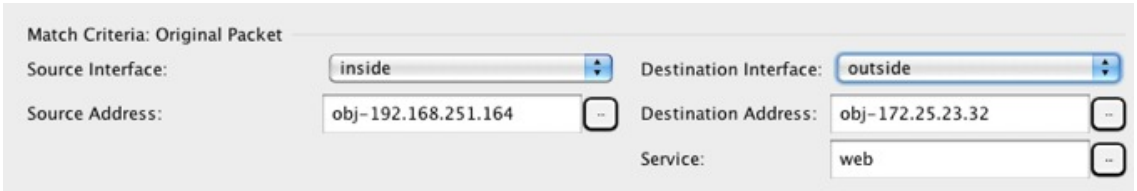
Step 6 (Optional.) Identify the destination service ports for service translation.

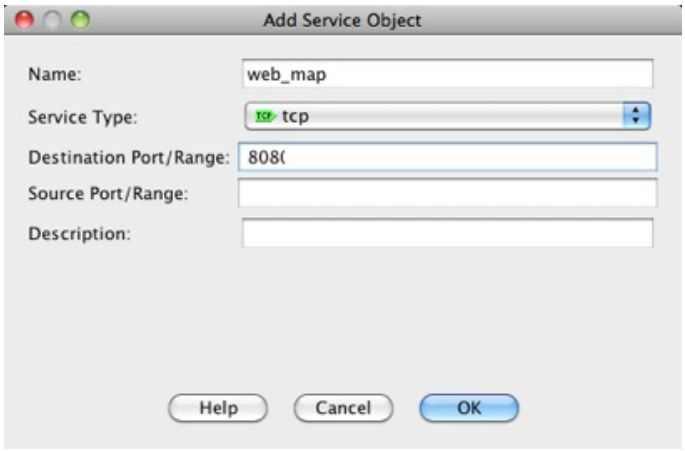
- Identify the original packet port (the *mapped destination port*). For **Match Criteria: Original Packet** > **Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Original Service dialog box.
- Identify the translated packet port (the *real destination port*). For **Action: Translated Packet** > **Service**, click the browse button and choose an existing service object that specifies TCP or UDP ports, or create a new object from the Browse Translated Service dialog box.

Dynamic NAT does not support port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

For example:









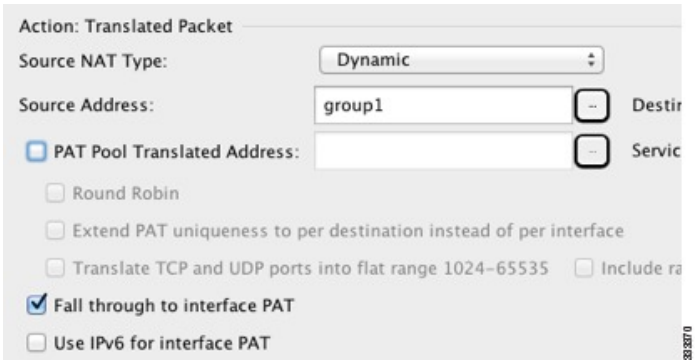
Step 7 (Optional.) For a PAT pool, configure the following options as needed:

- **Round Robin** —To assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.
- **Extend PAT uniqueness to per destination instead of per interface** (8.4(3) and later, not including 8.5(1) or 8.6(1).)—To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80.
- **Translate TCP or UDP ports into flat range (1024-65535)** (8.4(3) and later, not including 8.5(1) or 8.6(1).)—To use the 1024 to 65535 port range as a single flat range when allocating ports. When choosing the mapped port number for a translation, the ASA uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include range 1 to 1023** check box.
- **Include Reserved Ports (1 to 1023)**—Includes the reserved ports, 1-1023, in the range of ports that are available for address translation. If you do not specify this option, addresses are translated to ports in the 1024-65535 range only.
- **Enable Block Allocation** (9.5.1 and later)—Enables port block allocation. For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time. If you allocate a block of ports, subsequent connections from the host use new randomly selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Port blocks are allocated in the 1024-65535 range only. Port block allocation

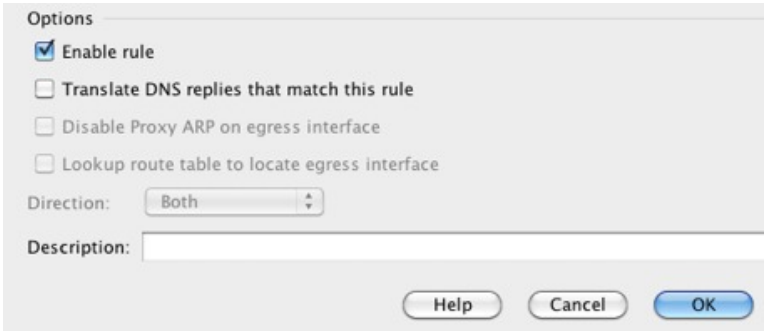
is compatible with round robin, but you cannot use it with extended PAT. You also cannot use interface PAT fallback.

Step 8 (Optional, when mapped interface is a non-bridge group member only.) To use the interface IP address as a backup method if the other mapped source addresses are already allocated, check the **Fall through to interface PAT** check box. To use the IPv6 interface address, also check the **Use IPv6 for interface PAT** check box.

The destination interface IP address is used. This option is only available if you configure a specific Destination Interface.



Step 9 (Optional.) Configure NAT options in the Options area.



- **Enable rule**—Enables this NAT rule. The rule is enabled by default.
- **Description**—Adds a description about the rule up to 200 characters in length.

Step 10 Click **OK**, then click **Apply**.

Configure PAT with Port Block Allocation

For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time (see RFC 6888). If you allocate a block of ports, subsequent connections from the host use new randomly-selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Blocks are freed when the last xlate that uses a port in the block is removed.

The main reason for allocating port blocks is reduced logging. The port block allocation is logged, connections are logged, but xlates created within the port block are not logged. On the other hand, this makes log analysis more difficult.

Port blocks are allocated in the 1024-65535 range only. There are separate blocks for TCP, UDP, and ICMP connections, and these blocks can overlap. Thus, if an application requires a low port number (1-1023), it might not work. For example, an application requesting port 22 (SSH) will get a mapped port within the range of 1024-65535 and within the block allocated to the host. You can create a separate NAT rule that does not use block allocation for applications that use low port numbers; for twice NAT, ensure the rule comes before the block allocation rule.

Before you begin

Usage notes for NAT rules:

- You can include the **Round Robin** option, but you cannot include the options for extending PAT uniqueness, or falling through to interface PAT. Other source/destination address and port information is also allowed.
- As with all NAT changes, if you replace an existing rule, you must clear xlates related to the replaced rule to have the new rule take effect. You can clear them explicitly or simply wait for them to time out. However, when operating in a cluster, you must reload the cluster so that the units can distribute blocks correctly.



Note If you are switching between a regular PAT and block allocation PAT rule, for object NAT, you must first delete the rule, then clear xlates. You can then create the new object NAT rule. Otherwise, you will see pat-port-block-state-mismatch drops in the **show asp drop** output.

- For a given PAT pool, you must specify (or not specify) block allocation for all rules that use the pool. You cannot allocate blocks in one rule and not in another. PAT pools that overlap also cannot mix block allocation settings. You also cannot overlap static NAT with port translation rules with the pool.

Procedure

Step 1 Select **Configuration** > **Firewall** > **Advanced** > **PAT Port Block Allocation** and configure the following settings:

- **Size of the block**—The number of ports in each block. The range is 32-4096. The default is 512.

  If you do not use the default, ensure that the size you choose divides evenly into 64,512 (the number of ports in the 1024-65535 range). Otherwise, there will be ports that cannot be used. For example, if you specify 100, there will be 12 unused ports.

- **Maximum block allocation per host**—The maximum number of blocks that can be allocated per host. The limit is per protocol, so a limit of 4 means at most 4 UDP blocks, 4 TCP blocks, and 4 ICMP blocks per host. The range is 1-8, the default is 4.
- **PBA Interim Logging**—If you enter a value, the system enables interim logging. By default, the system generates syslog messages during port block creation and deletion. If you enable interim logging, the

system generates the following message at the interval you specify. The messages report all active port blocks allocated at that time, including the protocol (ICMP, TCP, UDP) and source and destination interface and IP address, and the port block. You can specify an interval from 21600-604800 seconds (6 hours to 7 days).

%ASA-6-305017: Pba-interim-logging: Active *protocol* block of ports for translation from *real_interface*:*real_host_ip* to *mapped_interface*:*mapped_ip_address*/*start_port_num-end_port_num*

Step 2 Add NAT rules that use PAT pool block allocation.

a) Select **Configuration** > **Firewall** > **NAT Rules**.

b) Add or edit an object NAT or twice NAT rule.

c) Configure at least the following options:

- (Twice NAT.) Select the object that defines the source address in **Original Packet** > **Source Address**.
- **Type** = **Dynamic**.
- **Pat Pool Translated Address**. Select a network object that defines the pat pool network.
- **Enable Block Allocation**.

d) Click **OK**.

Configure Per-Session PAT or Multi-Session PAT (Version 9.0(1) and Higher)

By default, all TCP PAT traffic and all UDP DNS traffic uses per-session PAT. To use multi-session PAT for traffic, you can configure per-session PAT rules: a permit rule uses per-session PAT, and a deny rule uses multi-session PAT.

Per-session PAT improves the scalability of PAT and, for clustering, allows each member unit to own PAT connections; multi-session PAT connections have to be forwarded to and owned by the control unit. At the end of a per-session PAT session, the ASA sends a reset and immediately removes the xlate. This reset causes the end node to immediately release the connection, avoiding the TIME_WAIT state. Multi-session PAT, on the other hand, uses the PAT timeout, by default 30 seconds.

For "hit-and-run" traffic, such as HTTP or HTTPS, per-session PAT can dramatically increase the connection rate supported by one address. Without per-session PAT, the maximum connection rate for one address for an IP protocol is approximately 2000 per second. With per-session PAT, the connection rate for one address for an IP protocol is 65535/*average-lifetime*.

For traffic that can benefit from multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT by creating a per-session deny rule. However, if you also want to use per-session PAT for the UDP ports used by these protocols, you must create the permit rules for them.

Before you begin

By default, the following rules are installed:

- Permit TCP from any (IPv4 and IPv6) to any (IPv4 and IPv6).
- Permit UDP from any (IPv4 and IPv6) to the domain port.

These rules do not show up in the table.

You cannot remove these rules, and they always exist after any manually-created rules. Because rules are evaluated in order, you can override the default rules. For example, to completely negate these rules, you could add the following:

- Deny TCP from any (IPv4 and IPv6) to any (IPv4 and IPv6).
- Deny UDP from any (IPv4 and IPv6) to the domain port.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Advanced** > **Per-Session NAT Rules**.

Step 2 Do one of the following:

- Choose **Add > Add Per-Session NAT Rule**.
- Select a rule and click **Edit**.

Step 3 Configure the rule:

- **Action**—Click **Permit** or **Deny**. A permit rule uses per-session PAT; a deny rule uses multi-session PAT.
- Source—Specify the Source Address either by typing an address or clicking the **...** button to choose an object. For the service, select UDP or TCP. You can optionally specify a source port, although normally you only specify the destination port. Either type in UDP/*port* or TCP/*port*, or click the **...** button to select a common value or object.
- Destination—Specify the Destination Address either by typing an address or clicking the **...** button to choose an object. For the service, select UDP or TCP; this must match the source service. You can optionally specify a destination port. Either type in UDP/*port* or TCP/*port*, or click the **...** button to select a common value or object. You can use the operators != (not equal to), > (greater than), < (less than), or specify a range using a hyphen, for example, 100-200.

Step 4 Click **OK**, then click **Apply**.

Static NAT

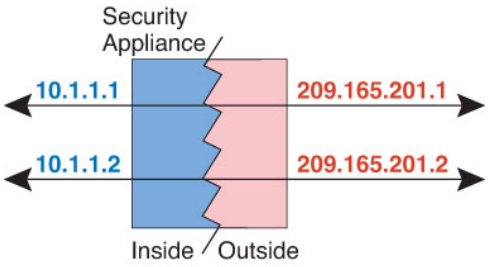
The following topics explain static NAT and how to implement it.

About static NAT

Static NAT creates a fixed translation of a real address to a mapped address. Because the mapped address is the same for each consecutive connection, static NAT allows bidirectional connection initiation, both to and from the host (if an access rule exists that allows it). With dynamic NAT and PAT, on the other hand, each host uses a different address or port for each subsequent translation, so bidirectional initiation is not supported.

The following figure shows a typical static NAT scenario. The translation is always active so both real and remote hosts can initiate connections.

Figure 21: Static NAT



Note You can disable bidirectionality if desired.

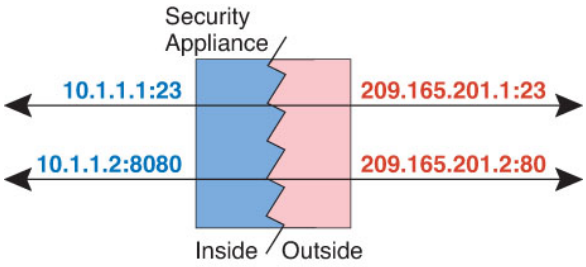
Static NAT with port translation

Static NAT with port translation lets you specify a real and mapped protocol and port.

When you specify the port with static NAT, you can choose to map the port and/or the IP address to the same value or to a different value.

The following figure shows a typical static NAT with port translation scenario showing both a port that is mapped to itself and a port that is mapped to a different value; the IP address is mapped to a different value in both cases. The translation is always active so both translated and remote hosts can initiate connections.

Figure 22: Typical Static NAT with Port Translation Scenario



Static NAT-with-port-translation rules limit access to the destination IP address for the specified port only. If you try to access the destination IP address on a different port not covered by a NAT rule, then the connection is blocked. In addition, for twice NAT, traffic that does not match the source IP address of the NAT rule will be dropped if it matches the destination IP address, regardless of the destination port. Therefore, you must add additional rules for all other traffic allowed to the destination IP address. For example, you can configure a static NAT rule for the IP address, without port specification, and place it after the port translation rule.

Note For applications that require application inspection for secondary channels (for example, FTP and VoIP), NAT automatically translates the secondary ports.

Following are some other uses of static NAT with port translation.

Static NAT with Identity Port Translation

You can simplify external access to internal resources. For example, if you have three separate servers that provide services on different ports (such as FTP, HTTP, and SMTP), you can give external users a single IP address to access those services. You can then configure static NAT with identity port translation to map the single external IP address to the correct IP addresses of the real servers based on the port they are trying to access. You do not need to change the port, because the servers are using the standard ones (21, 80, and 25 respectively). For details on how to configure this example, see Single Address for FTP, HTTP, and SMTP (Static NAT-with-Port-Translation), on page 224.

Static NAT with Port Translation for Non-Standard Ports

You can also use static NAT with port translation to translate a well-known port to a non-standard port or vice versa. For example, if inside web servers use port 8080, you can allow outside users to connect to port 80, and then undo translation to the original port 8080. Similarly, to provide extra security, you can tell web users to connect to non-standard port 6785, and then undo translation to port 80.

Static Interface NAT with Port Translation

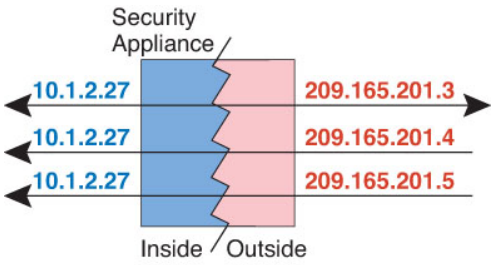
You can configure static NAT to map a real address to an interface address/port combination. For example, if you want to redirect Telnet access for the device's outside interface to an inside host, then you can map the inside host IP address/port 23 to the outside interface address/port 23.

One-to-many static NAT

Typically, you configure static NAT with a one-to-one mapping. However, in some cases, you might want to configure a single real address to several mapped addresses (one-to-many). When you configure one-to-many static NAT, when the real host initiates traffic, it always uses the first mapped address. However, for traffic initiated to the host, you can initiate traffic to any of the mapped addresses, and they will be untranslated to the single real address.

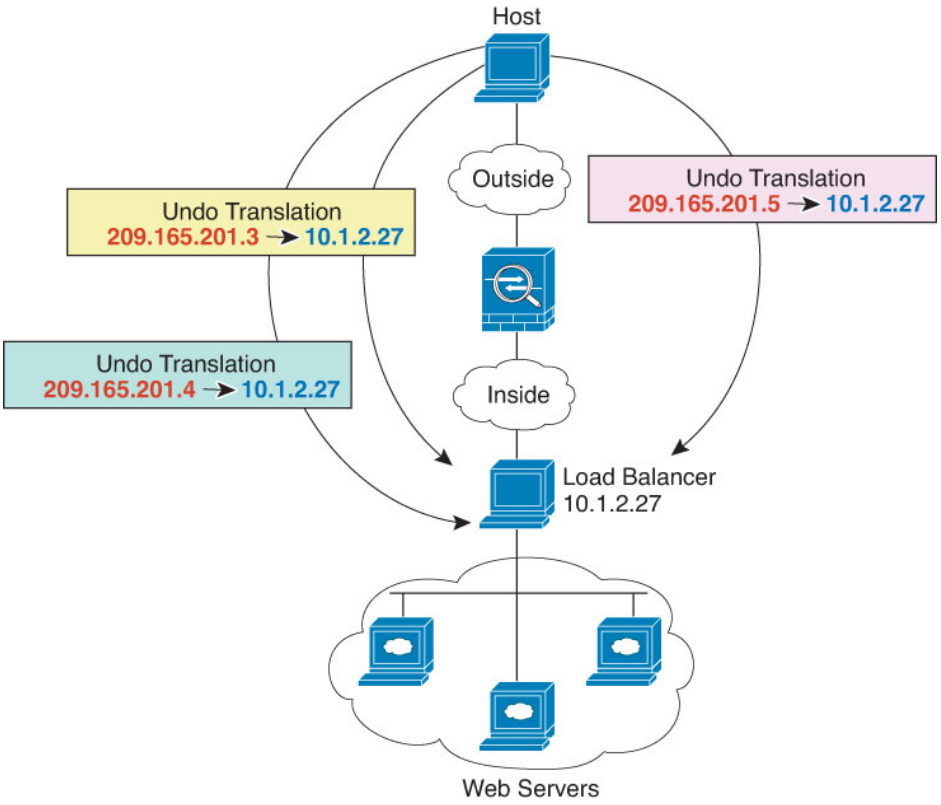
The following figure shows a typical one-to-many static NAT scenario. Because initiation by the real host always uses the first mapped address, the translation of real host IP/first mapped IP is technically the only bidirectional translation.

Figure 23: One-to-Many Static NAT



For example, you have a load balancer at 10.1.2.27. Depending on the URL requested, it redirects traffic to the correct web server. For details on how to configure this example, see Inside Load Balancer with Multiple Mapped Addresses (Static NAT, One-to-Many), on page 222.

Figure 24: One-to-Many Static NAT Example



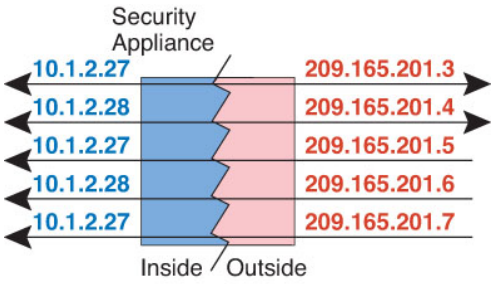
Other mapping scenarios (not recommended)

NAT has the flexibility to allow any kind of static mapping scenario: one-to-one, one-to-many, but also few-to-many, many-to-few, and many-to-one mappings. We recommend using only one-to-one or one-to-many mappings. These other mapping options might result in unintended consequences.

Functionally, few-to-many is the same as one-to-many; but because the configuration is more complicated and the actual mappings may not be obvious at a glance, we recommend creating a one-to-many configuration for each real address that requires it. For example, for a few-to-many scenario, the few real addresses are mapped to the many mapped addresses in order (A to 1, B to 2, C to 3). When all real addresses are mapped, the next mapped address is mapped to the first real address, and so on until all mapped addresses are mapped (A to 4, B to 5, C to 6). This results in multiple mapped addresses for each real address. Just like a one-to-many configuration, only the first mappings are bidirectional; subsequent mappings allow traffic to be initiated *to* the real host, but all traffic *from* the real host uses only the first mapped address for the source.

The following figure shows a typical few-to-many static NAT scenario.

Figure 25: Few-to-Many Static NAT

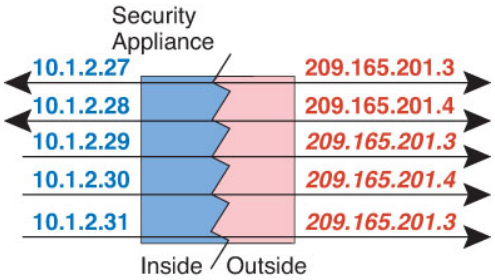


For a many-to-few or many-to-one configuration, where you have more real addresses than mapped addresses, you run out of mapped addresses before you run out of real addresses. Only the mappings between the lowest real IP addresses and the mapped pool result in bidirectional initiation. The remaining higher real addresses can initiate traffic, but traffic cannot be initiated to them (returning traffic for a connection is directed to the correct real address because of the unique 5-tuple (source IP, destination IP, source port, destination port, protocol) for the connection).

Note Many-to-few or many-to-one NAT is not PAT. If two real hosts use the same source port number and go to the same outside server and the same TCP destination port, and both hosts are translated to the same IP address, then both connections will be reset because of an address conflict (the 5-tuple is not unique).

The following figure shows a typical many-to-few static NAT scenario.

Figure 26: Many-to-Few Static NAT



Instead of using a static rule this way, we suggest that you create a one-to-one rule for the traffic that needs bidirectional initiation, and then create a dynamic rule for the rest of your addresses.

Configure Static Network Object NAT or Static NAT-with-Port-Translation

This section describes how to configure a static NAT rule using network object NAT.

Procedure

Step 1 Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration** > **Firewall** > **NAT Rules**, then click **Add** > **Add Network Object NAT Rule**.
- To add NAT to an existing network object, choose **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups**, and then edit a network object.

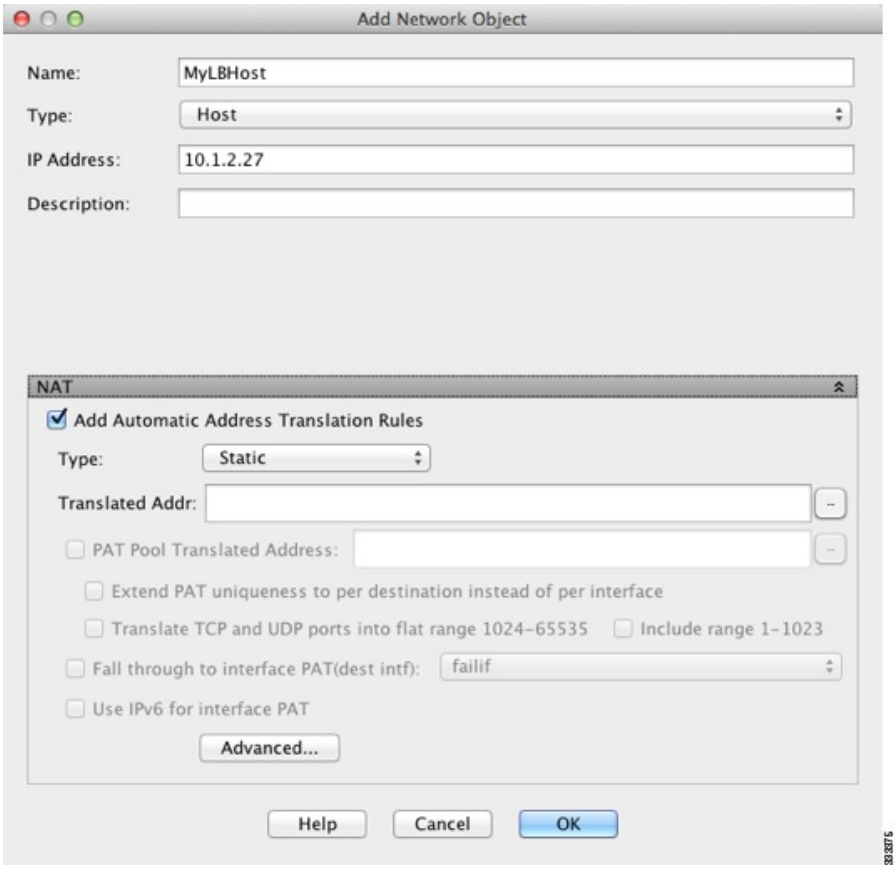
Step 2 For a new object, enter values for the following fields:

- **Name**—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.
- **Type**—Host, Network, or Range.
- **IP Addresses**—IPv4 or IPv6 addresses, a single address for a host, a starting and ending address for a range, and for subnet, either an IPv4 network address and mask (for example, 10.100.10.0 255.255.255.0) or IPv6 address and prefix length (for example, 2001:DB8:0:CD30::/60).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

Step 4 Check the **Add Automatic Translation Rules** check box.

Step 5 From the Type drop-down list, choose **Static**.



Step 6 In the Translated Addr. field, specify the mapped IP address as one of the following. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses. For more information, see Static NAT, on page 189.

- Type a host IP address. This provides a one-to-one mapping for host objects. For subnet objects, the same netmask is used for the inline host address, and you get one-to-one translations for addresses in the mapped inline host's subnet. For range objects, the mapped address includes the same number of hosts that are in the range object, starting with the mapped host address. For example, if the real address is defined as a range from 10.1.1.1 through 10.1.1.6, and you specify 172.20.1.1 as the mapped address, then the mapped range will include 172.20.1.1 through 172.20.1.6. For NAT46 or NAT66 translations, this can be an IPv6 network address.
- Click the browse button and select a network object (or create a new one). To do a one-to-one mapping for a range of IP addresses, select an object that contains a range with the same number of addresses.
- (For static NAT-with-port-translation only.) Type an interface name or click the browse button, and choose an interface from the Browse Translated Addr dialog box. You cannot select a bridge group member interface.



To use the IPv6 interface address, you must also check the **Use IPv6 for interface PAT** check box. Be sure to also click **Advanced** and configure a service port translation. (You cannot specify an interface in transparent mode.)

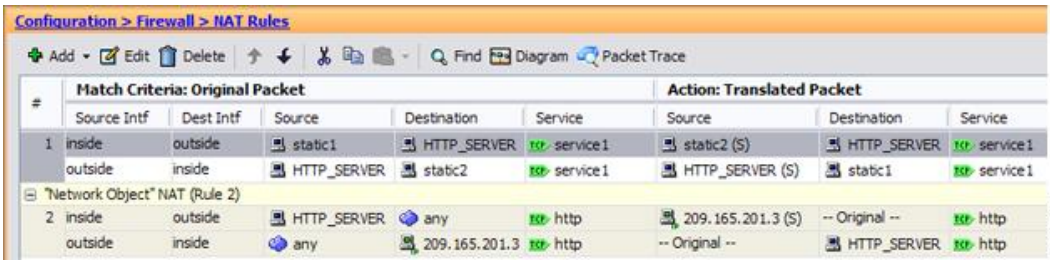
Step 7 (Optional.) For NAT46, check **Use one-to-one address translation**. For NAT 46, specify one-to-one to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this keyword.

Step 8 (Optional) Click **Advanced**, configure the following options in the Advanced NAT Settings dialog box, and click **OK**.

- **Translate DNS replies for rule**—Translates the IP address in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). See Rewriting DNS queries and responses using NAT, on page 262 for more information.
- **Disable Proxy ARP on egress interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. For information on the conditions which might require the disabling of proxy ARP, see Mapped Addresses and Routing, on page 244.
- (Required for bridge group member interfaces.) **Interface**—Specifies the real interface (**Source**) and the mapped interface (**Destination**) where this NAT rule applies. By default, the rule applies to all interfaces except for bridge group members.
- **Service**—Configures static NAT-with-port-translation. Choose the protocol, then enter the real port and the mapped port. You can use port numbers or a well-known port name such as http.

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

Because static rules are bidirectional (allowing initiation to and from the real host), the NAT Rules table shows two rows for each static rule, one for each direction.



Configure Static Twice NAT or Static NAT-with-Port-Translation

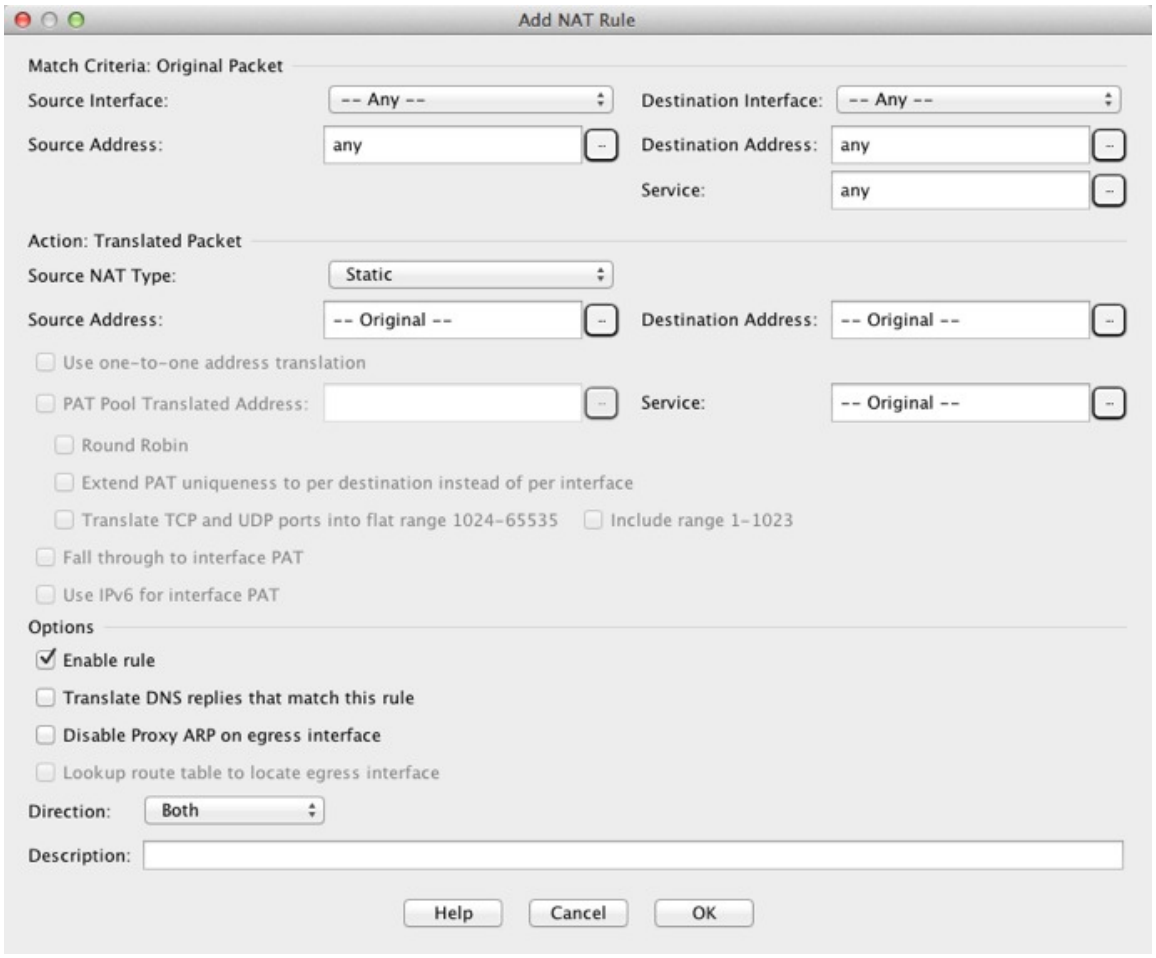
This section describes how to configure a static NAT rule using twice NAT.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**, and then do one of the following:

- Click **Add**, or **Add > Add NAT Rule Before Network Object NAT Rules**.
- Click **Add > Add NAT Rule After Network Object NAT Rules**.
- Select a twice NAT rule and click **Edit**.

The Add NAT Rule dialog box appears.

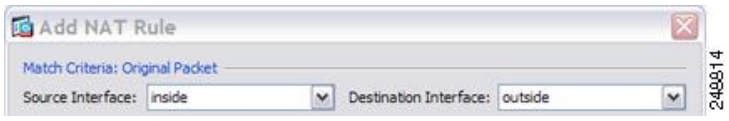


Step 2 (Required for bridge group member interfaces.) Set the source and destination interfaces.

In routed mode, the default is any interface for both source and destination. You can select specific interfaces for one or both options. However, you must select the interface when writing a rule for bridge group member interfaces; "any" does not include these interfaces.

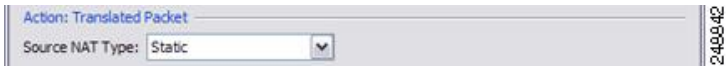
a) From the **Match Criteria: Original Packet > Source Interface** drop-down list, choose the source interface.

b) From the **Match Criteria: Original Packet > Destination Interface** drop-down list, choose the destination interface.

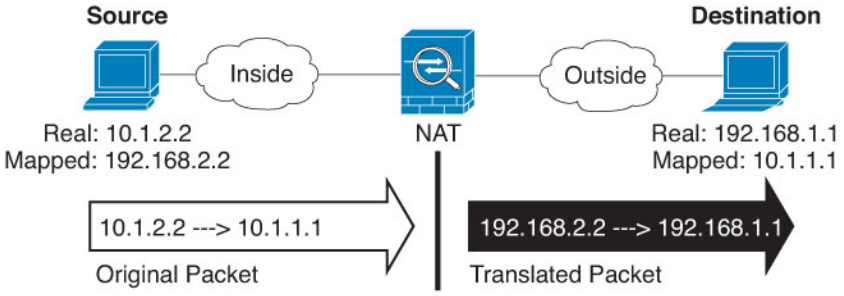


Step 3 Choose **Static** from the **Action: Translated Packet** > **Source NAT Type** drop-down list. Static is the default setting.

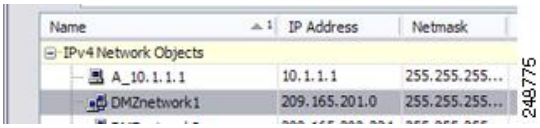
This setting only applies to the source address; the destination translation is always static.



Step 4 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.



a) For **Match Criteria: Original Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**, but do not use this option except for identity NAT.



b) (Optional) For **Match Criteria: Original Packet > Destination Address**, click the browse button and choose an existing network object, group, or interface, or create a new object or group from the Browse Original Destination Address dialog box.

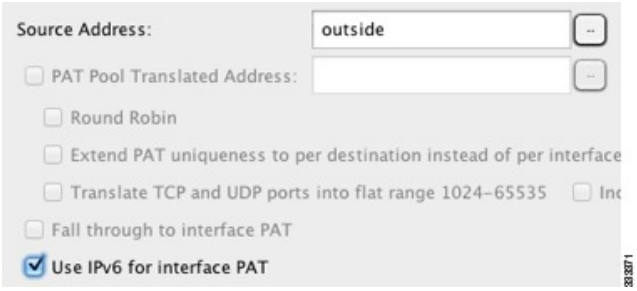
Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 152.

Step 5 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired.

a) For **Action: Translated Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated Source Address dialog box.

For static NAT, the mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired.

For static interface NAT with port translation, you can specify the interface instead of a network object/group for the mapped address. If you want to use the IPv6 address of the interface, check the **Use IPv6 for interface PAT** check box. You cannot select a bridge group member interface.



For more information, see Static NAT with port translation, on page 190. See Guidelines for NAT, on page 155 for information about disallowed mapped IP addresses.

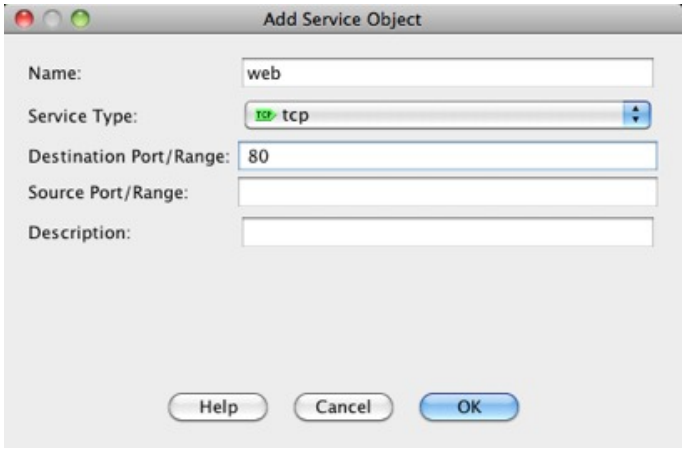
b) (Optional.) For **Action: Translated Packet > Destination Address**, click the browse button and choose an existing network object or group, or create a new object or group from the Browse Translated Destination Address dialog box.

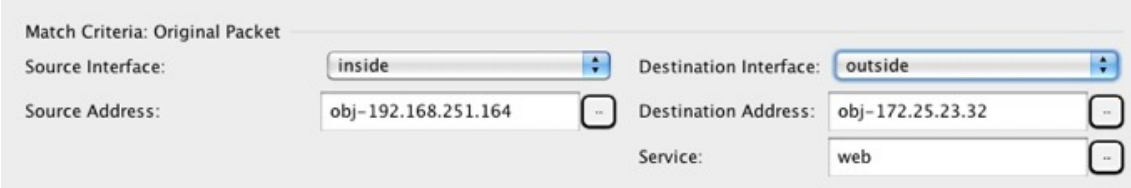
Step 6 (Optional.) Identify the source or destination service ports for service translation.

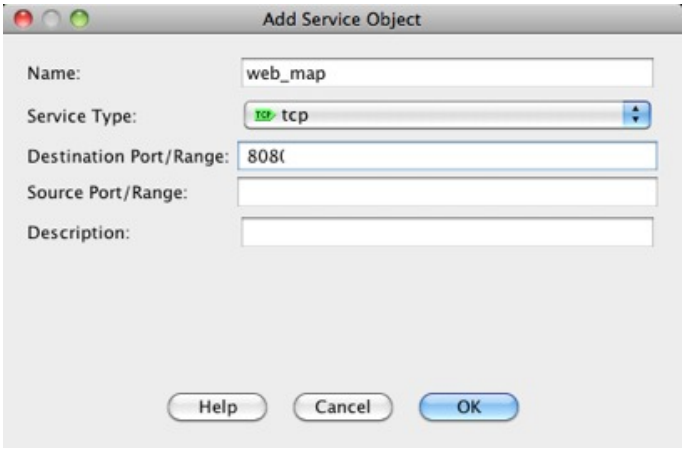
- Identify the original packet source or destination port (the *real source port* **or** the *mapped destination port*). For **Match Criteria: Original Packet > Service**, click the browse button and choose an existing service object that specifies ports, or create a new object from the Browse Original Service dialog box.
- Identify the translated packet source or destination port (the *mapped source port* **or** the *real destination port*). For **Action: Translated Packet > Service**, click the browse button and choose an existing service object that specifies ports, or create a new object from the Browse Translated Service dialog box.

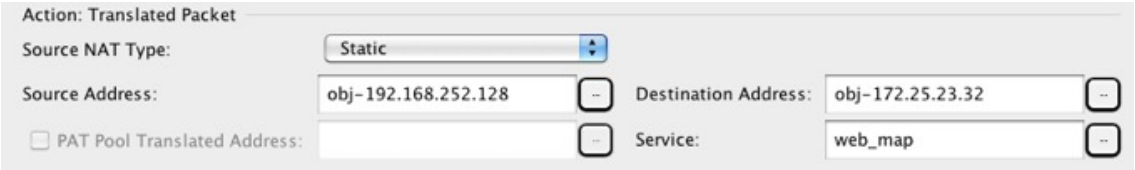
A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both the real and mapped service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real destination port. When translating a port, be sure the protocols in the real and mapped service objects are identical (for example, both TCP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

For example:



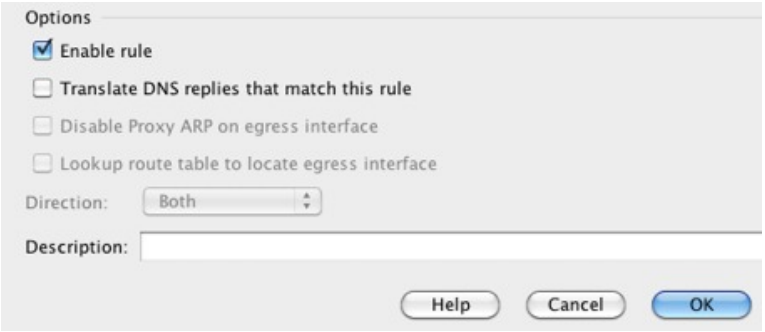






Step 7 (Optional.) For NAT46, check the **Use one-to-one address translation** check box. For NAT46, specify one-to-one to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this keyword.

Step 8 (Optional.) Configure NAT options in the Options area.



- **Enable rule** —Enables this NAT rule. The rule is enabled by default.
- (For a source-only rule.) **Translate DNS replies that match this rule**—Rewrites the DNS A record in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). You cannot configure DNS modification if you configure a destination address. See Rewriting DNS queries and responses using NAT, on page 262 for more information.

- **Disable Proxy ARP on egress interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. See Mapped Addresses and Routing, on page 244 for more information.
- **Direction**—To make the rule unidirectional, choose **Unidirectional**. The default is Both. Making the rule unidirectional prevents destination addresses from initiating connections to the real addresses.
- **Description**—Adds a description about the rule up to 200 characters in length.

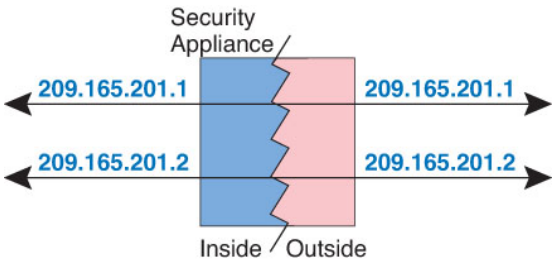
Step 9 Click **OK**, then click **Apply**.

Identity NAT

You might have a NAT configuration in which you need to translate an IP address to itself. For example, if you create a broad rule that applies NAT to every network, but want to exclude one network from NAT, you can create a static NAT rule to translate that network to itself. Identity NAT is necessary for remote access VPN, where you need to exempt the client traffic from NAT.

The following figure shows a typical identity NAT scenario.

Figure 27: Identity NAT



The following topics explain how to configure identity NAT.

Configure Identity Network Object NAT

This section describes how to configure an identity NAT rule using network object NAT.

Procedure

Step 1 Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration** > **Firewall** > **NAT Rules**, then click **Add** > **Add Network Object NAT Rule**.
- To add NAT to an existing network object, choose **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups**, and then edit a network object.

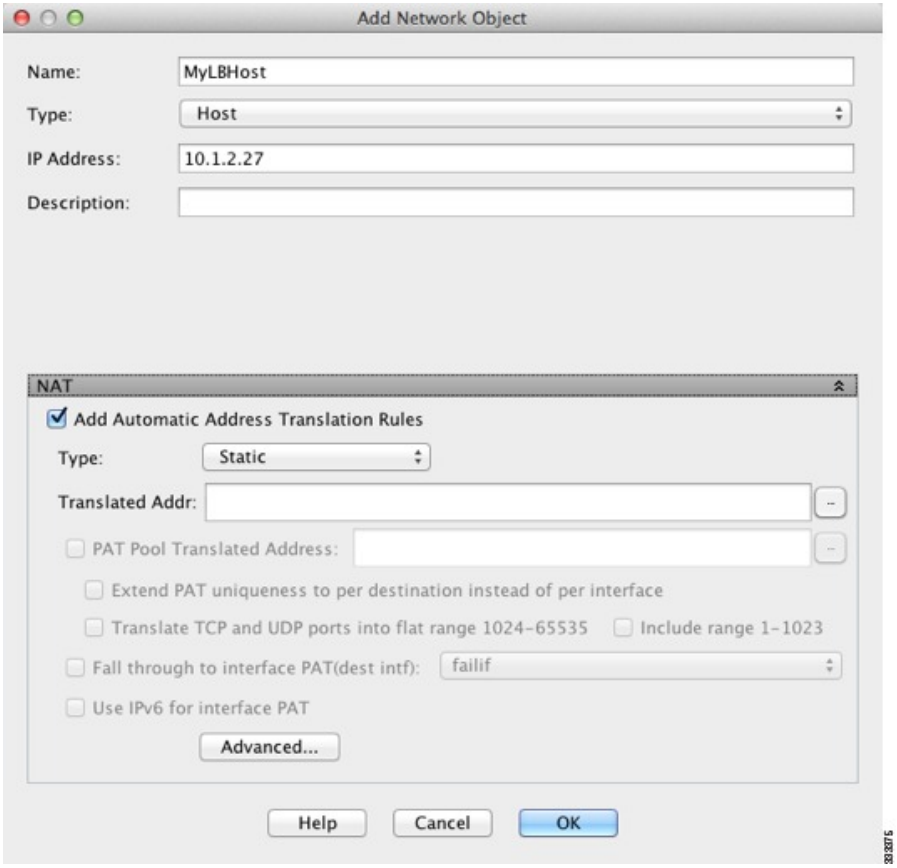
Step 2 For a new object, enter values for the following fields:

- **Name**—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.
- **Type**—Host, Network, or Range.
- **IP Addresses**—IPv4 or IPv6 addresses, a single address for a host, a starting and ending address for a range, and for subnet, either an IPv4 network address and mask (for example, 10.100.10.0 255.255.255.0) or IPv6 address and prefix length (for example, 2001:DB8:0:CD30::/60).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

Step 4 Check the **Add Automatic Translation Rules** check box.

Step 5 From the Type drop-down list, choose **Static**.



Step 6 In the Translated Addr. field, do one of the following:

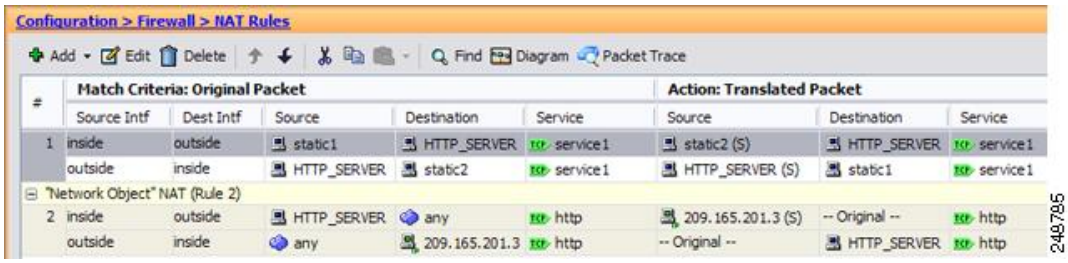
- For host objects, enter the same address. For range objects, enter the first address in the real range (the same number of addresses in the range will be used). For subnet objects, enter any address within the real subnet (all addresses in the subnet will be used).
- Click the browse button and select a network object (or create a new one). Use this option when configuring identity NAT for a range of addresses.

Step 7 (Optional) Click **Advanced**, configure the following options in the Advanced NAT Settings dialog box, and click **OK**.

- **Translate DNS replies for rule**—Do not configure this option for identity NAT.
- **Disable Proxy ARP on egress interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. For information on the conditions which might require the disabling of proxy ARP, see Mapped Addresses and Routing, on page 244.
- (Routed mode; interfaces specified.) **Lookup route table to locate egress interface**—Determines the egress interface using a route lookup instead of using the interface specified in the NAT command. See Determining the Egress Interface, on page 246 for more information.
- (Required for bridge group member interfaces.) **Interface**—Specifies the real interface (**Source**) and the mapped interface (**Destination**) where this NAT rule applies. By default, the rule applies to all interfaces except for bridge group members.
- **Service**—Do not configure this option for identity NAT.

Step 8 Click **OK**, and then **Apply**.

Because static rules are bidirectional (allowing initiation to and from the real host), the NAT Rules table shows two rows for each static rule, one for each direction, unless you select the route lookup option.



Configure Identity Twice NAT

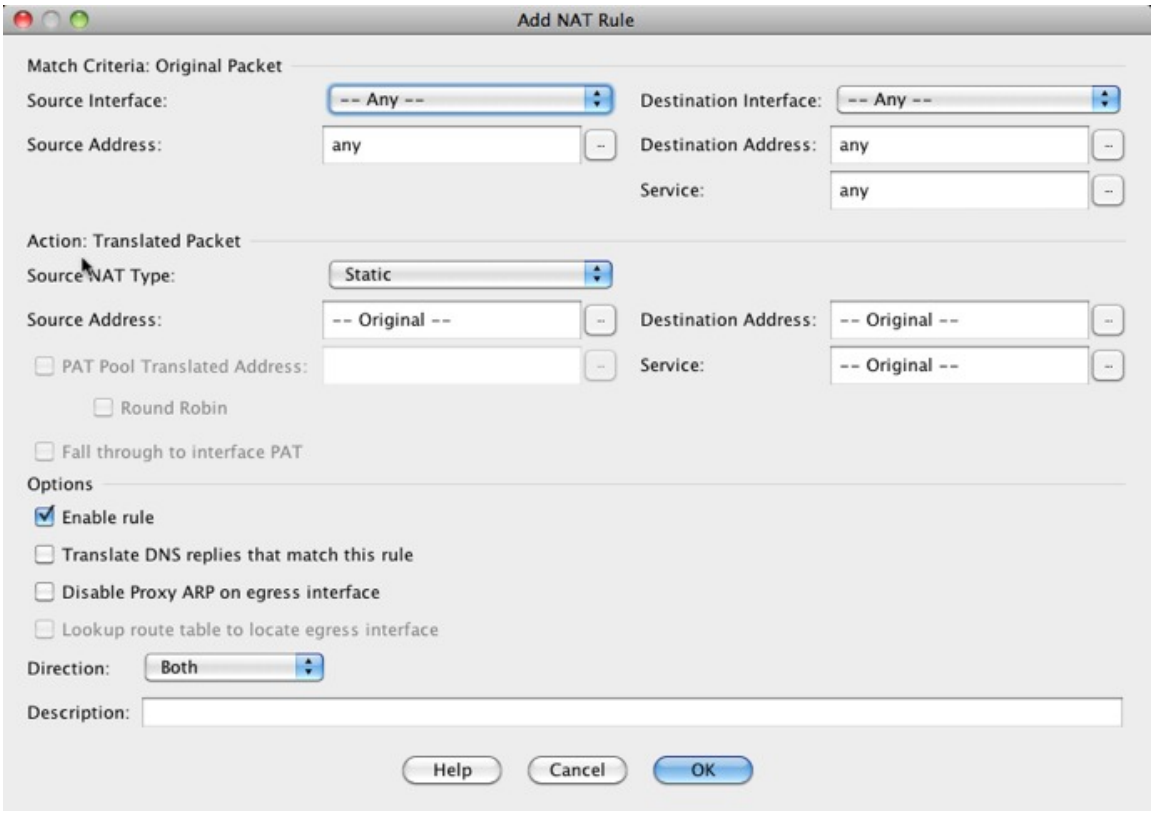
This section describes how to configure an identity NAT rule using twice NAT.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**, and then do one of the following:

- Click **Add**, or **Add > Add NAT Rule Before Network Object NAT Rules**.
- Click **Add > Add NAT Rule After Network Object NAT Rules**.
- Select a twice NAT rule and click **Edit**.

The Add NAT Rule dialog box appears.

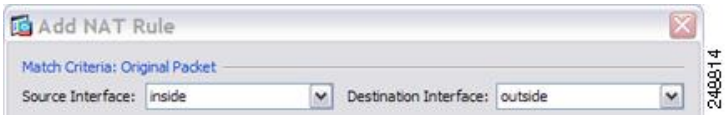


Step 2 (Required for bridge group member interfaces.) Set the source and destination interfaces.

In routed mode, the default is any interface for both source and destination. You can select specific interfaces for one or both options. However, you must select the interface when writing a rule for bridge group member interfaces; "any" does not include these interfaces.

a) From the **Match Criteria: Original Packet > Source Interface** drop-down list, choose the source interface.

b) From the **Match Criteria: Original Packet > Destination Interface** drop-down list, choose the destination interface.



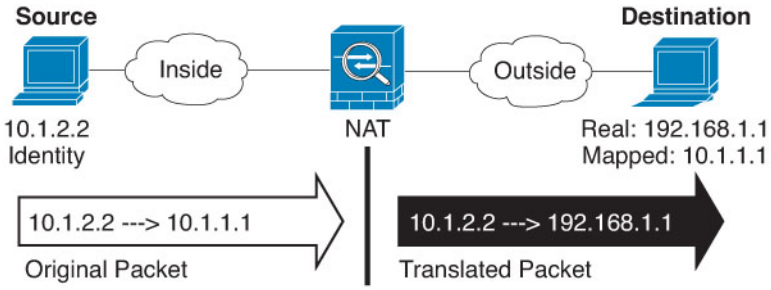
Step 3 Choose **Static** from the **Action: Translated Packet** > **Source NAT Type** drop-down list. Static is the default setting.

This setting only applies to the source address; the destination translation is always static.

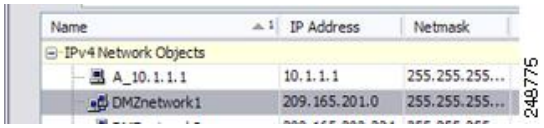


Step 4 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following

figure for an example of the original packet vs. the translated packet where you perform identity NAT on the inside host but translate the outside host.



a) For **Match Criteria: Original Packet > Source Address**, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**; only use this option when also setting the mapped address to **any**.



b) (Optional.) For **Match Criteria: Original Packet > Destination Address**, click the browse button and choose an existing network object, group, or interface, or create a new object or group from the Browse Original Destination Address dialog box.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Comparing Network object NAT and Twice NAT, on page 152.

For static interface NAT with port translation only, choose an interface. If you specify an interface, be sure to also configure a service translation. For more information, see Static NAT with port translation, on page 190.

Step 5 Identify the translated packet addresses; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*).

a) For **Action: Translated Packet > Source Address**, click the browse button and choose the same network object or group from the Browse Translated Source Address dialog box that you chose for the real source address. Use **any** if you specified **any** for the real address.

b) For **Match Criteria: Translated Packet > Destination Address**, click the browse button and choose an existing network object or group, or create a new object or group from the Browse Translated Destination Address dialog box.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

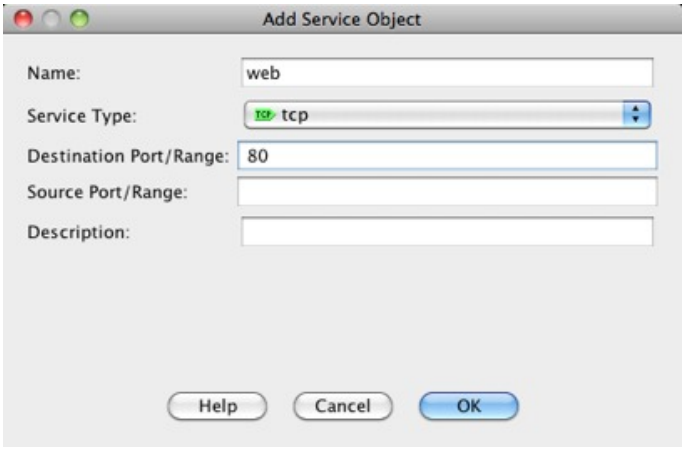
If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see Static NAT, on page 189. See Guidelines for NAT, on page 155 for information about disallowed mapped IP addresses.

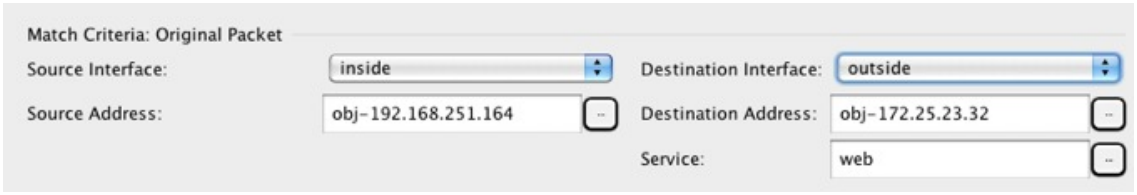
Step 6 (Optional.) Identify the source or destination service ports for service translation.

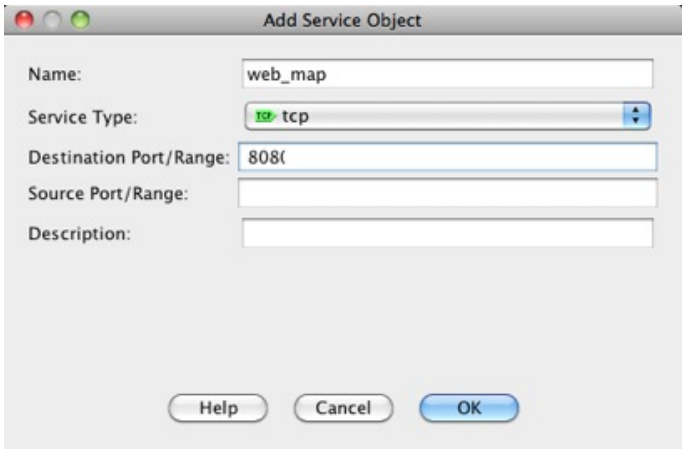
- Identify the original packet source or destination port (the *real source port* **or** the *mapped destination port*). For **Match Criteria: Original Packet > Service**, click the browse button and choose an existing service object that specifies ports, or create a new object from the Browse Original Service dialog box.
- Identify the translated packet source or destination port (the *mapped source port* **or** the *real destination port*). For **Action: Translated Packet > Service**, click the browse button and choose an existing service object that specifies ports, or create a new object from the Browse Translated Service dialog box.

A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both the real and mapped service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real destination port. When translating a port, be sure the protocols in the real and mapped service objects are identical (for example, both TCP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

For example:

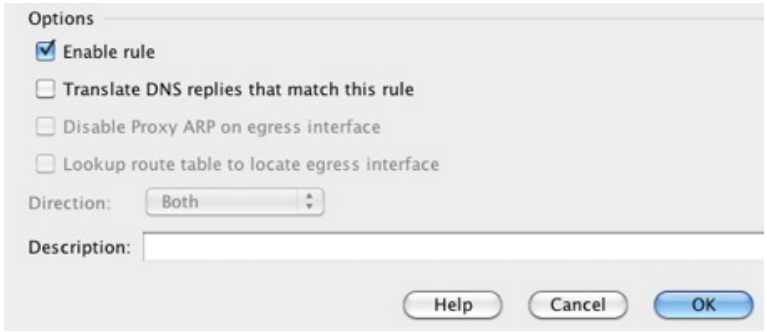








Step 7 (Optional) Configure NAT options in the Options area.



- **Enable rule** —Enables this NAT rule. The rule is enabled by default.
- (For a source-only rule.) **Translate DNS replies that match this rule**—Although this option is available if you do not configure a destination address, it is not applicable to identity NAT because you are translating the address to itself, so the DNS reply does not need modification.
- **Disable Proxy ARP on egress interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. See Mapped Addresses and Routing, on page 244 for more information.
- (Routed mode; interfaces specified.) **Lookup route table to locate egress interface**—Determines the egress interface using a route lookup instead of using the interface specified in the NAT command. See Determining the Egress Interface, on page 246 for more information.
- **Direction**—To make the rule unidirectional, choose **Unidirectional**. The default is Both. Making the rule unidirectional prevents traffic from initiating connections to the real addresses. You might want to use this setting for testing purposes.
- **Description**—Adds a description about the rule up to 200 characters in length.

Step 8 Click **OK**, then click **Apply**.

Monitoring NAT

You can view NAT related graphs from the following pages:

- **Monitoring > Properties > Connection Graphs > Xlates**—Select the Xlate Utilization graph to view the in-use and most-used xlates. This is equivalent to the **show xlate** command.
- **Monitoring > Properties > Connection Graphs > Perfmon**—Select the Xlate Perfmon graph to see NAT performance information. This is equivalent to the xlate information from the **show perfmon** command.

History for NAT

Feature Name	Platform Releases	Description
Network Object NAT	8.3(1)	Configures NAT for a network object IP address(es). We introduced or modified the following screens: Configuration > Firewall > NAT Rules Configuration > Firewall > Objects > Network Objects/Groups
Twice NAT	8.3(1)	Twice NAT lets you identify both the source and destination address in a single rule. We modified the following screen: Configuration > Firewall > NAT Rules.

Feature Name	Platform Releases	Description
Identity NAT configurable proxy ARP and route lookup	8.4(2)/8.5(1)	In earlier releases for identity NAT, proxy ARP was disabled, and a route lookup was always used to determine the egress interface. You could not configure these settings. In 8.4(2) and later, the default behavior for identity NAT was changed to match the behavior of other static NAT configurations: proxy ARP is enabled, and the NAT configuration determines the egress interface (if specified) by default. You can leave these settings as is, or you can enable or disable them discretely. Note that you can now also disable proxy ARP for regular static NAT. For pre-8.3 configurations, the migration of NAT exempt rules (the **nat 0 access-list** command) to 8.4(2) and later now includes the following keywords to disable proxy ARP and to use a route lookup: **no-proxy-arp** and **route-lookup**. The **unidirectional** keyword that was used for migrating to 8.3(2) and 8.4(1) is no longer used for migration. When upgrading to 8.4(2) from 8.3(1), 8.3(2), and 8.4(1), all identity NAT configurations will now include the **no-proxy-arp** and **route-lookup** keywords, to maintain existing functionality. The **unidirectional** keyword is removed. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit Network Object > Advanced NAT Settings; Configuration > Firewall > NAT Rules > Add/Edit NAT Rule.
PAT pool and round robin address assignment	8.4(2)/8.5(1)	You can now specify a pool of PAT addresses instead of a single address. You can also optionally enable round-robin assignment of PAT addresses instead of first using all ports on a PAT address before using the next address in the pool. These features help prevent a large number of connections from a single PAT address from appearing to be part of a DoS attack and makes configuration of large numbers of PAT addresses easy. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit Network Object; Configuration > Firewall > NAT Rules > Add/Edit NAT Rule.
Round robin PAT pool allocation uses the same IP address for existing hosts	8.4(3)	When using a PAT pool with round robin allocation, if a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. We did not modify any screens. *This feature is not available in 8.5(1) or 8.6(1).*

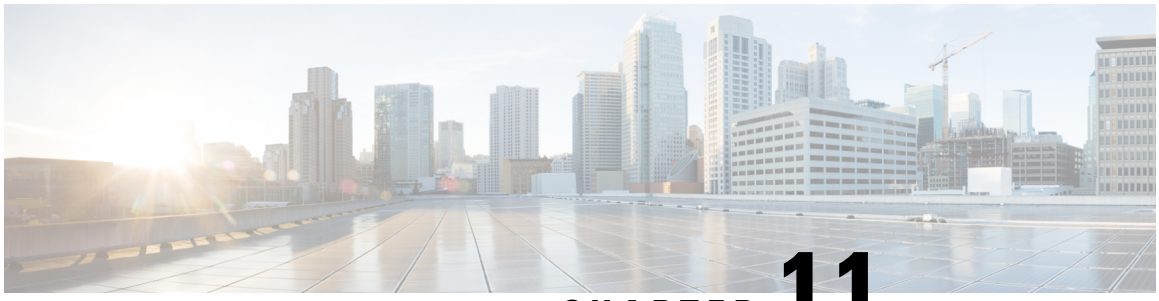
Feature Name	Platform Releases	Description
Flat range of PAT ports for a PAT pool	8.4(3)	If available, the real source port number is used for the mapped port. However, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 0 to 511, 512 to 1023, and 1024 to 65535. Therefore, ports below 1024 have only a small PAT pool. If you have a lot of traffic that uses the lower port ranges, when using a PAT pool, you can now specify a flat range of ports to be used instead of the three unequal-sized tiers: either 1024 to 65535, or 1 to 65535. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit Network Object; Configuration > Firewall > NAT Rules > Add/Edit NAT Rule. *This feature is not available in 8.5(1) or 8.6(1).*
Extended PAT for a PAT pool	8.4(3)	Each PAT IP address allows up to 65535 ports. If 65535 ports do not provide enough translations, you can now enable extended PAT for a PAT pool. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit Network Object; Configuration > Firewall > NAT Rules > Add/Edit NAT Rule. *This feature is not available in 8.5(1) or 8.6(1).*

Feature Name	Platform Releases	Description
Automatic NAT rules to translate a VPN peer's local IP address back to the peer's real IP address	8.4(3)	In rare situations, you might want to use a VPN peer's real IP address on the inside network instead of an assigned local IP address. Normally with VPN, the peer is given an assigned local IP address to access the inside network. However, you might want to translate the local IP address back to the peer's real public IP address if, for example, your inside servers and network security is based on the peer's real IP address. You can enable this feature on one interface per tunnel group. Object NAT rules are dynamically added and deleted when the VPN session is established or disconnected. You can view the rules using the **show nat** command. Because of routing issues, we do not recommend using this feature unless you know you need it; contact Cisco TAC to confirm feature compatibility with your network. See the following limitations: • Only supports Cisco IPsec and AnyConnect Client. • Return traffic to the public IP addresses must be routed back to the ASA so the NAT policy and VPN policy can be applied. • Does not support load-balancing (because of routing issues). • Does not support roaming (public IP changing). ASDM does not support this command; enter the command using the Command Line Tool.
NAT support for IPv6	9.0(1)	NAT now supports IPv6 traffic, as well as translating between IPv4 and IPv6. Translating between IPv4 and IPv6 is not supported in transparent mode. We modified the following screen: Configuration > Firewall > Objects > Network Objects/Group; Configuration > Firewall > NAT Rules.
NAT support for reverse DNS lookups	9.0(1)	NAT now supports translation of the DNS PTR record for reverse DNS lookups when using IPv4 NAT, IPv6 NAT, and NAT64 with DNS inspection enabled for the NAT rule.

Feature Name	Platform Releases	Description
Per-session PAT	9.0(1)	The per-session PAT feature improves the scalability of PAT and, for clustering, allows each member unit to own PAT connections; multi-session PAT connections have to be forwarded to and owned by the control unit. At the end of a per-session PAT session, the ASA sends a reset and immediately removes the xlate. This reset causes the end node to immediately release the connection, avoiding the TIME_WAIT state. Multi-session PAT, on the other hand, uses the PAT timeout, by default 30 seconds. For "hit-and-run" traffic, such as HTTP or HTTPS, the per-session feature can dramatically increase the connection rate supported by one address. Without the per-session feature, the maximum connection rate for one address for an IP protocol is approximately 2000 per second. With the per-session feature, the connection rate for one address for an IP protocol is 65535/*average-lifetime*. By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate. For traffic that requires multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT by creating a per-session deny rule. We introduced the following screen: Configuration > Firewall > Advanced > Per-Session NAT Rules.
Transactional Commit Model on NAT Rule Engine	9.3(1)	When enabled, a NAT rule update is applied after the rule compilation is completed; without affecting the rule matching performance. We added NAT to the following screen: Configuration > Device Management > Advanced > Rule Engine.
Carrier Grade NAT enhancements	9.5(1)	For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time (see RFC 6888). We added the following command: **Configuration** > **Firewall** > **Advanced** > **PAT Port Block Allocation**. We added **Enable Block Allocation** the object NAT and twice NAT dialog boxes.
NAT support for SCTP	9.5(2)	You can now specify SCTP ports in static network object NAT rules. Using SCTP in static twice NAT is not recommended. Dynamic NAT/PAT does not support SCTP. We modified the following screen: **Configuration** > **Firewall** > **NAT** add/edit static network object NAT rule, **Advanced NAT Settings** dialog box.

Feature Name	Platform Releases	Description
Interim logging for NAT port block allocation.	9.12(1)	When you enable port block allocation for NAT, the system generates syslog messages during port block creation and deletion. If you enable interim logging, the system generates message 305017 at the interval you specify. The messages report all active port blocks allocated at that time, including the protocol (ICMP, TCP, UDP) and source and destination interface and IP address, and the port block. We modified the following screen: **Configuration** > **Firewall** > **Advanced** > **PAT Port Block Allocation**.
Changes to PAT address allocation in clustering. The PAT pool **flat** option is now enabled by default and it is not configurable.	9.15(1)	The way PAT addresses are distributed to the members of a cluster is changed. Previously, addresses were distributed to the members of the cluster, so your PAT pool would need a minimum of one address per cluster member. Now, the control unit instead divides each PAT pool address into equal-sized port blocks and distributes them across cluster members. Each member has port blocks for the same PAT addresses. Thus, you can reduce the size of the PAT pool, even to as few as one IP address, depending on the amount of connections you typically need to PAT. Port blocks are allocated in 512-port blocks from the 1024-65535 range. You can optionally included the reserved ports, 1-1023, in this block allocation when you configure PAT pool rules. For example, in a 4-node cluster, each node gets 32 blocks with which it will be able to handle 16384 connections per PAT pool IP address compared to a single node handling all 65535 connections per PAT pool IP address. As part of this change, PAT pools for all systems, whether standalone or operating in a cluster, now use a flat port range of 1023 - 65535. Previously, you could optionally use a flat range by including the **flat** keyword in a PAT pool rule. The **flat** keyword is no longer supported: the PAT pool is now always flat. The **include-reserve** keyword, which was previously a sub-keyword to **flat**, is now an independent keyword within the PAT pool configuration. With this option, you can include the 1 - 1023 port range within the PAT pool. Note that if you configure port block allocation (the **block-allocation** PAT pool option), your block allocation size is used rather than the default 512-port block. In addition, you cannot configure extended PAT for a PAT pool for systems in a cluster. New/Modified screens: NAT PAT Pool configuration.

Feature Name	Platform Releases	Description
New Section 0 for system-defined NAT rules.	9.16(1)	A new Section 0 has been added to the NAT rule table. This section is exclusively for the use of the system. Any NAT rules that the system needs for normal functioning are added to this section, and these rules take priority over any rules you create. Previously, system-defined rules were added to Section 1, and user-defined rules could interfere with proper system functioning. You cannot add, edit, or delete Section 0 rules, but you will see them in **show nat detail** command output.



CHAPTER 11

NAT Examples and Reference

The following topics provide examples for configuring NAT, plus information on advanced configuration and troubleshooting.

- Examples for Network Object NAT, on page 215
- Examples for Twice NAT, on page 228
- NAT in routed and transparent mode, on page 241
- Routing NAT Packets, on page 243
- NAT for VPN, on page 247
- Translating IPv6 networks, on page 252
- Rewriting DNS queries and responses using NAT, on page 262

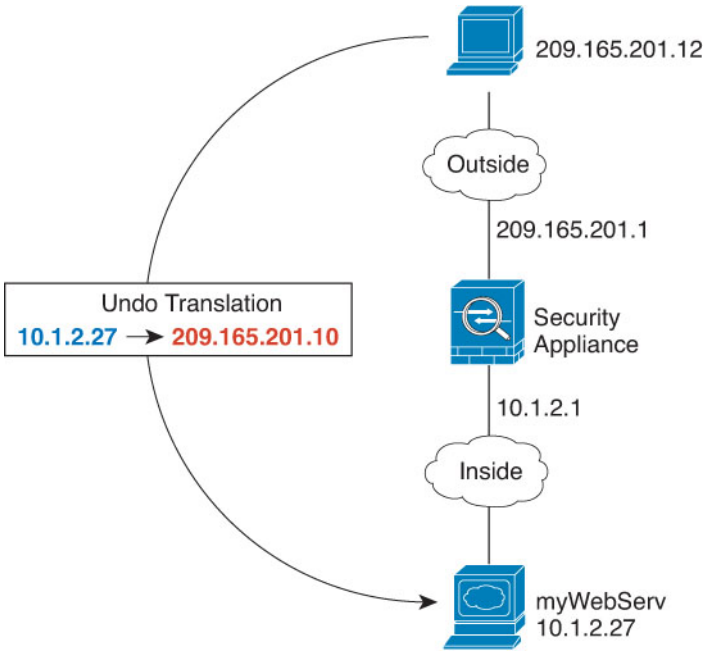
Examples for Network Object NAT

Following are some configuration examples for network object NAT.

Providing Access to an Inside Web Server (Static NAT)

The following example performs static NAT for an inside web server. The real address is on a private network, so a public address is required. Static NAT is necessary so hosts can initiate traffic to the web server at a fixed address.

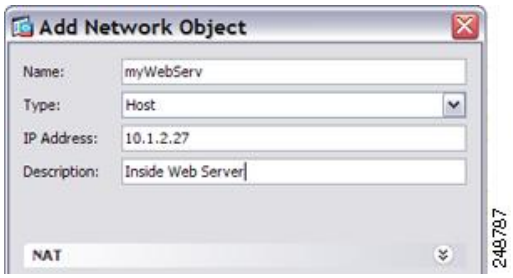
Figure 28: Static NAT for an Inside Web Server



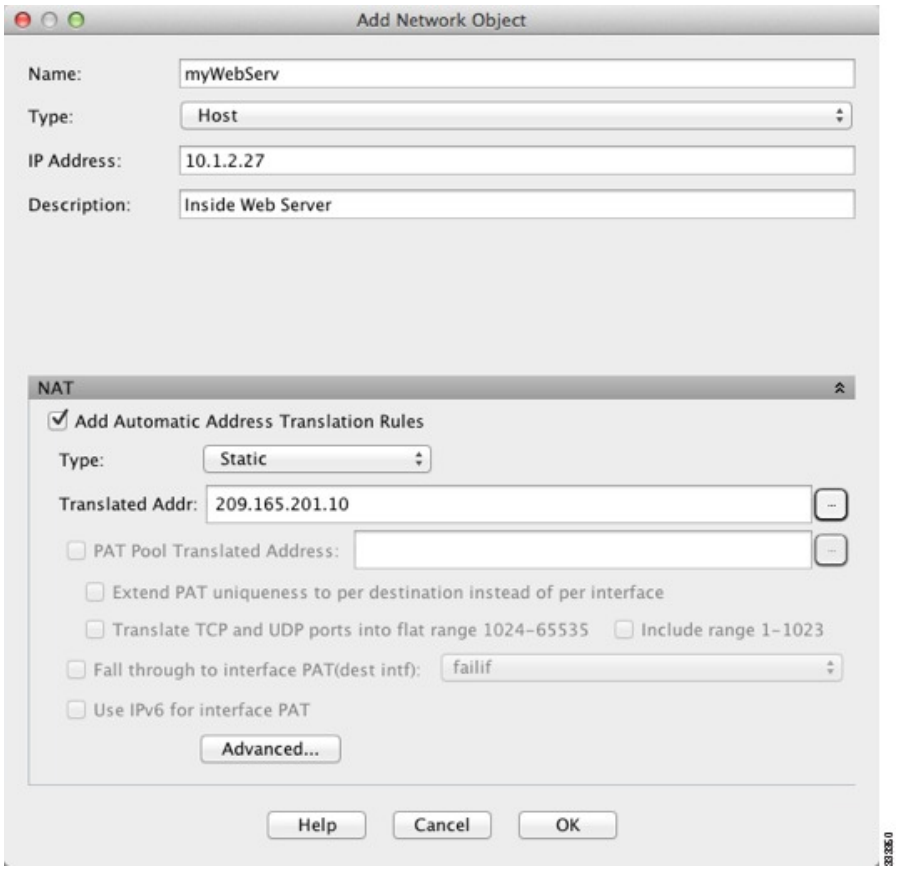
Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT**.

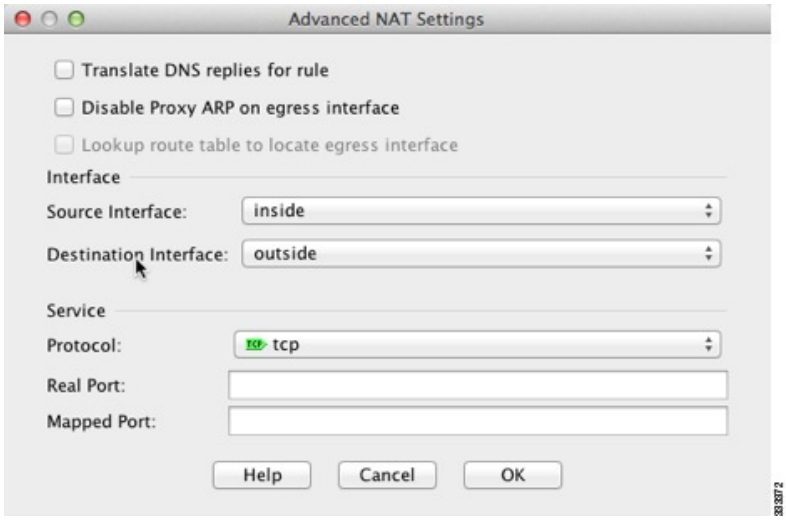
Step 2 Choose **Add** > **Network Object NAT Rule**, name the new network object and define the web server host address.



Step 3 Configure static NAT for the object.



Step 4 Click **Advanced** and configure the real and mapped interfaces.

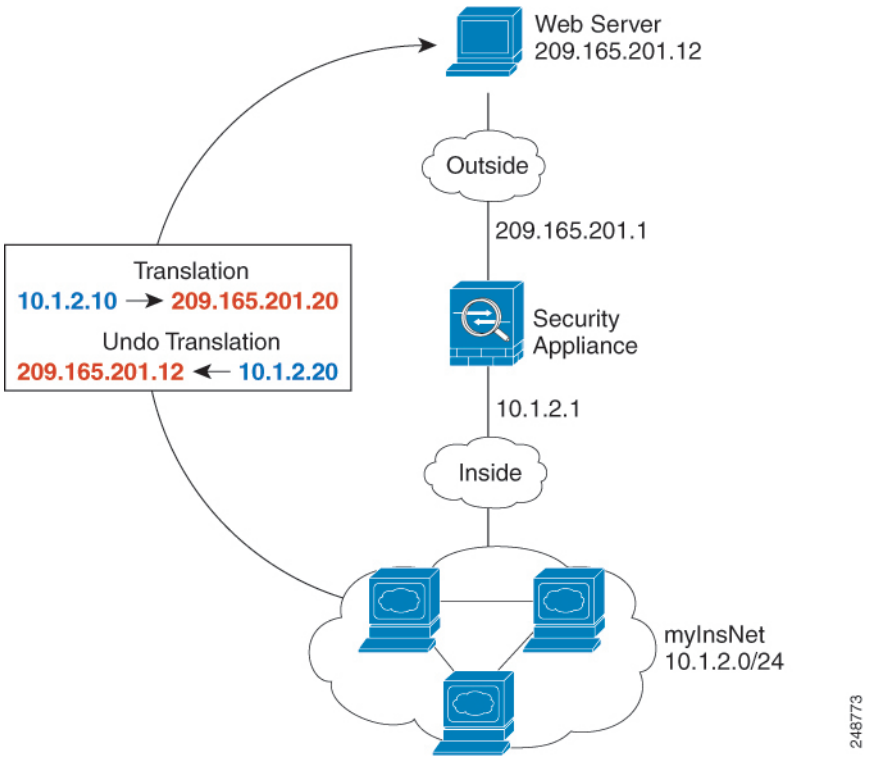


Step 5 Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

NAT for Inside Hosts (Dynamic NAT) and NAT for an Outside Web Server (Static NAT)

The following example configures dynamic NAT for inside users on a private network when they access the outside. Also, when inside users connect to an outside web server, that web server address is translated to an address that appears to be on the inside network.

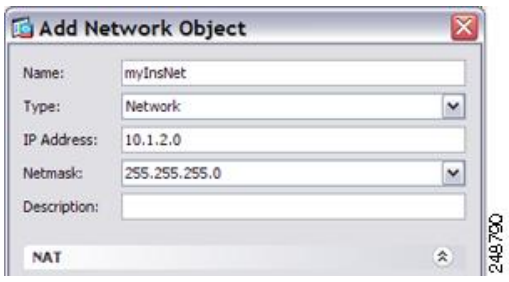
Figure 29: Dynamic NAT for Inside, Static NAT for Outside Web Server



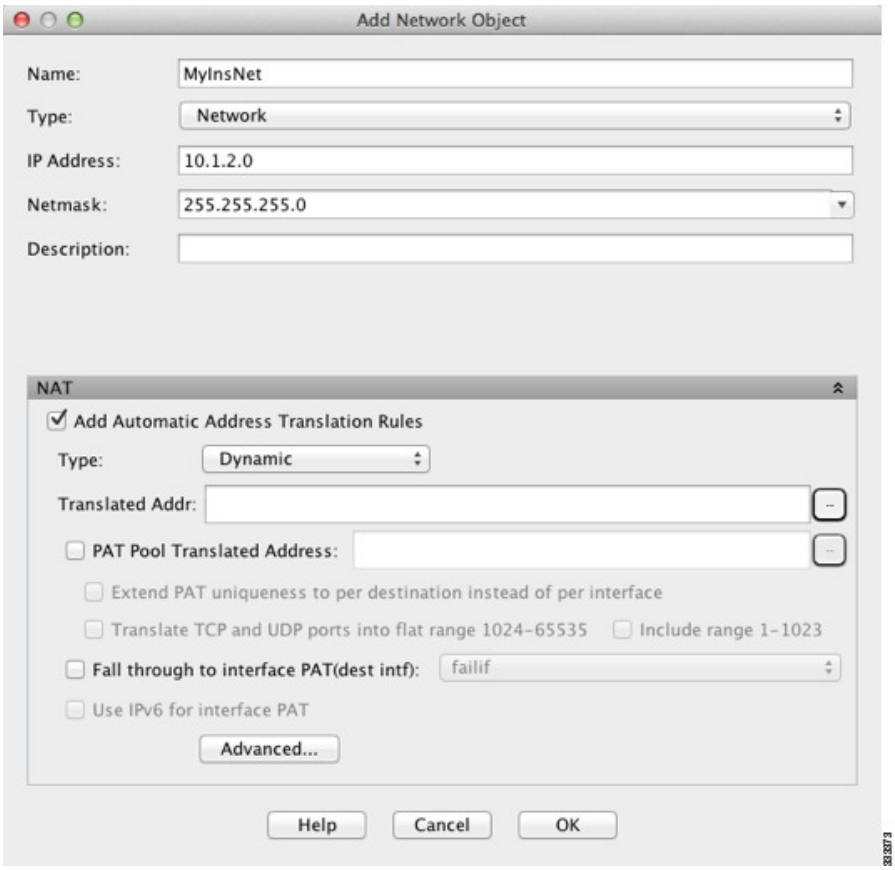
Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT**.

Step 2 Choose **Add** > **Network Object NAT Rule**, name the new network object and define the inside network.

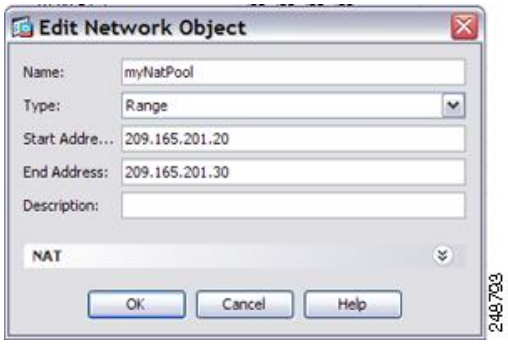


Step 3 Enable dynamic NAT for the inside network.

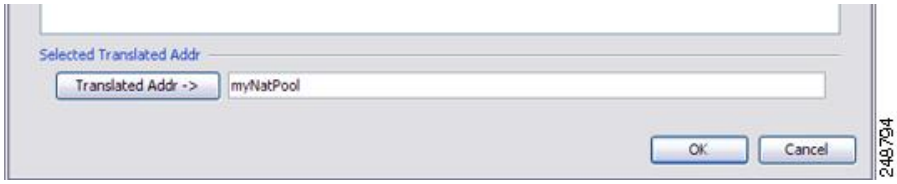


Step 4 For the Translated Addr field, add a new network object for the dynamic NAT pool to which you want to translate the inside addresses by clicking the browse button.

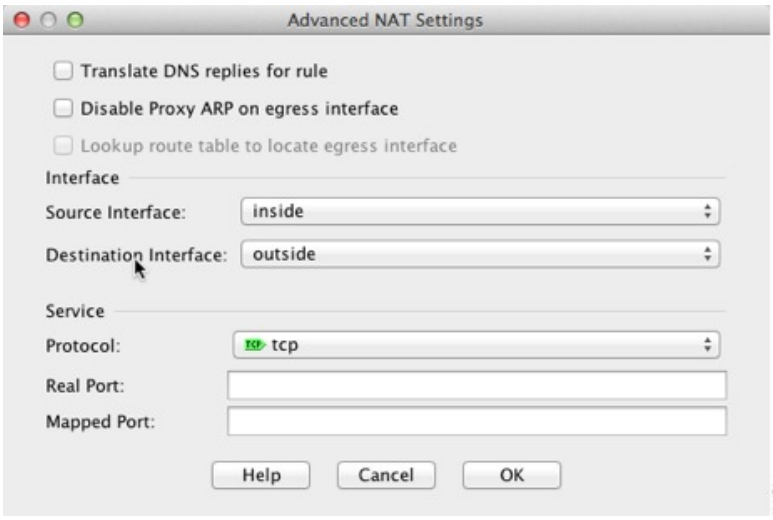
a) Choose **Add > Network Object**, name the new object, define the range of addresses in the NAT pool, and click **OK**.



b) Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

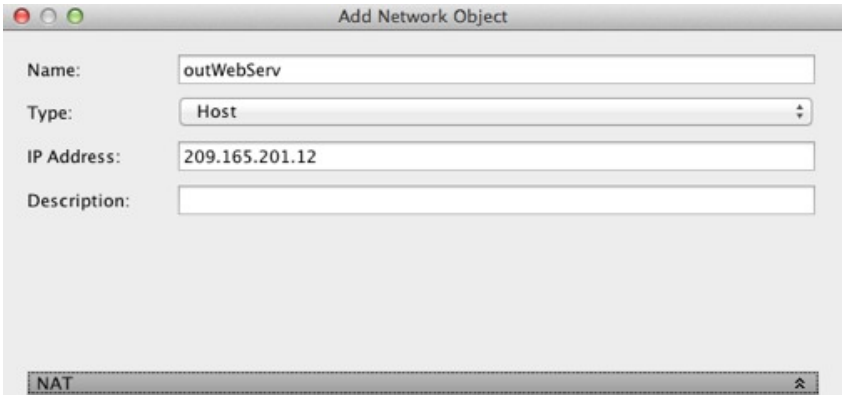


Step 5 Click **Advanced** and configure the real and mapped interfaces.

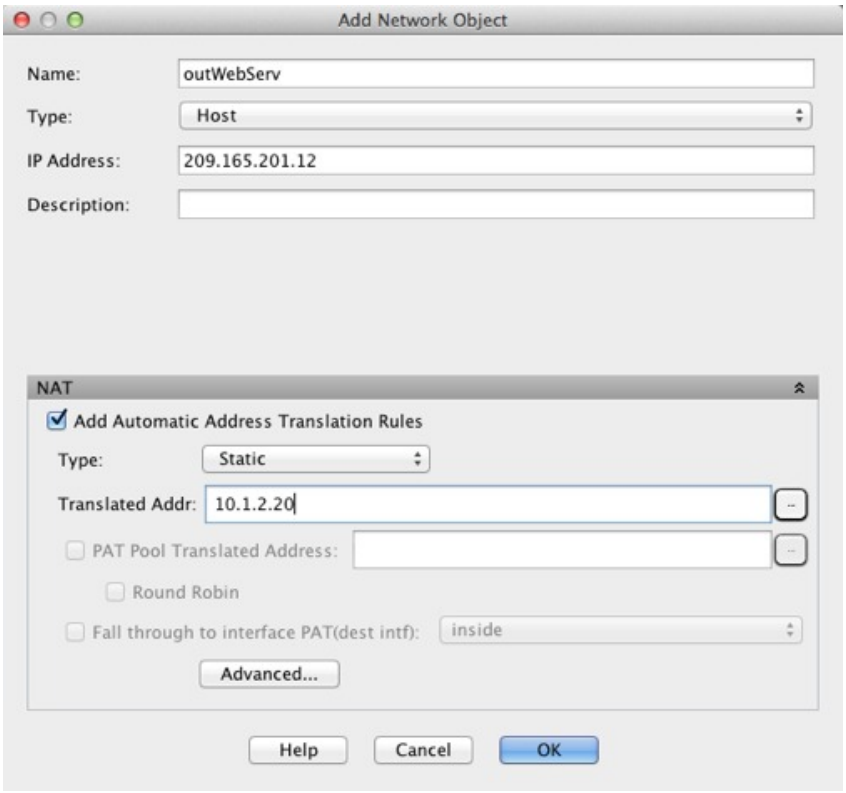


Step 6 Click **OK** to return to the Edit Network Object dialog box, click then click **OK** again to return to the NAT Rules table.

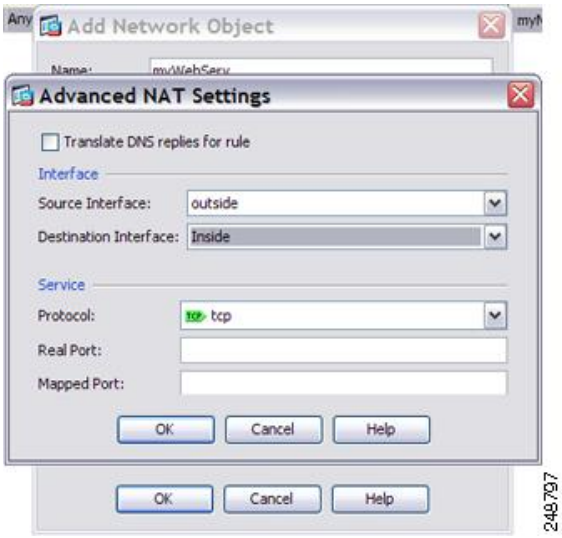
Step 7 Choose **Add** > **Network Object NAT Rule** and create an object for the outside web server.



Step 8 Configure static NAT for the web server.



Step 9 Click **Advanced** and configure the real and mapped interfaces.

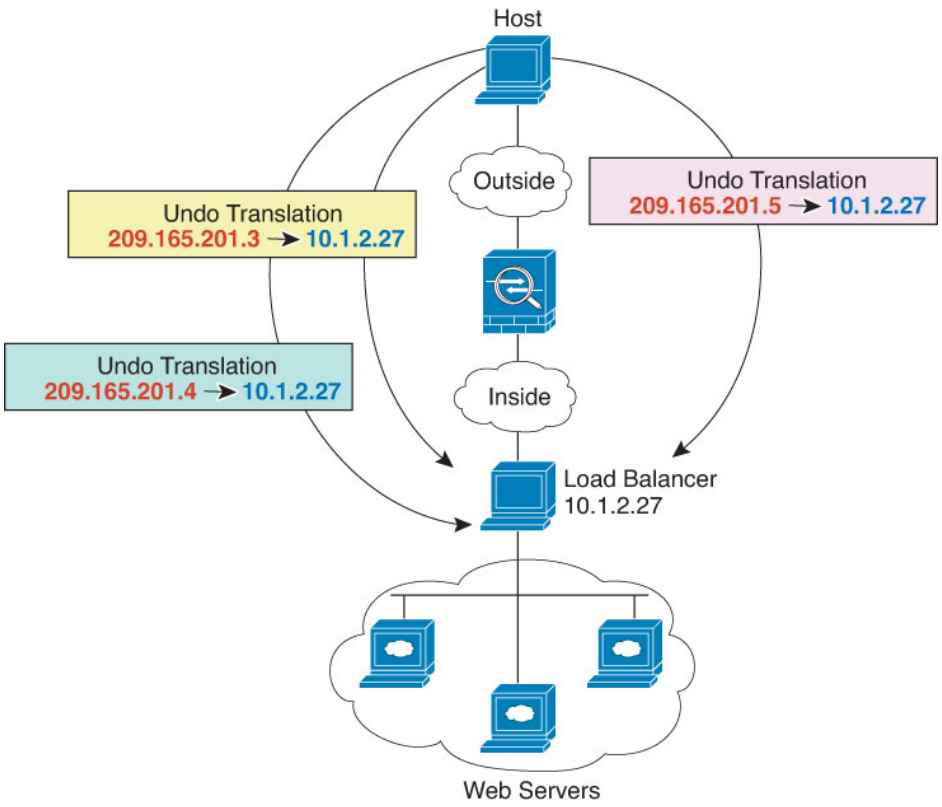


Step 10 Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

Inside Load Balancer with Multiple Mapped Addresses (Static NAT, One-to-Many)

The following example shows an inside load balancer that is translated to multiple IP addresses. When an outside host accesses one of the mapped IP addresses, it is untranslated to the single load balancer address. Depending on the URL requested, it redirects traffic to the correct web server.

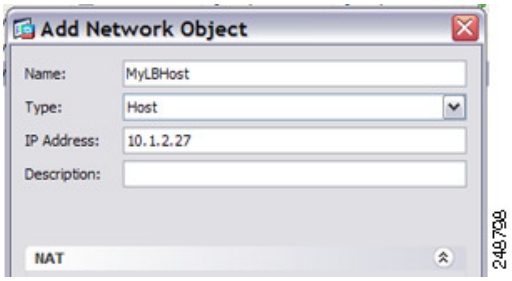
Figure 30: Static NAT with One-to-Many for an Inside Load Balancer



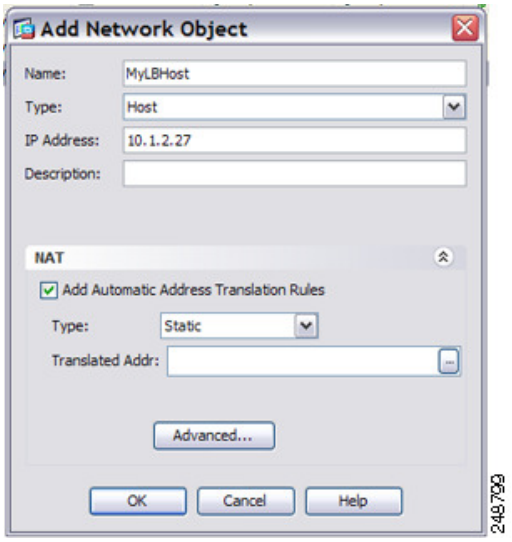
Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT**.

Step 2 Choose **Add** > **Network Object NAT Rule**, name the new network object and define the load balancer address.

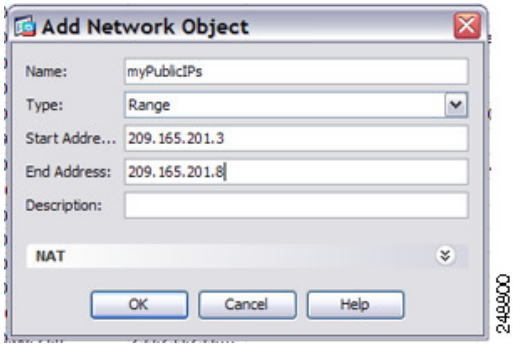


Step 3 Enable static NAT for the load balancer:

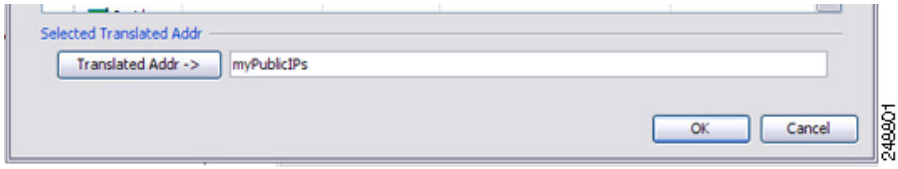


Step 4 For the Translated Addr field, add a new network object for the static NAT group of addresses to which you want to translate the load balancer address by clicking the browse button.

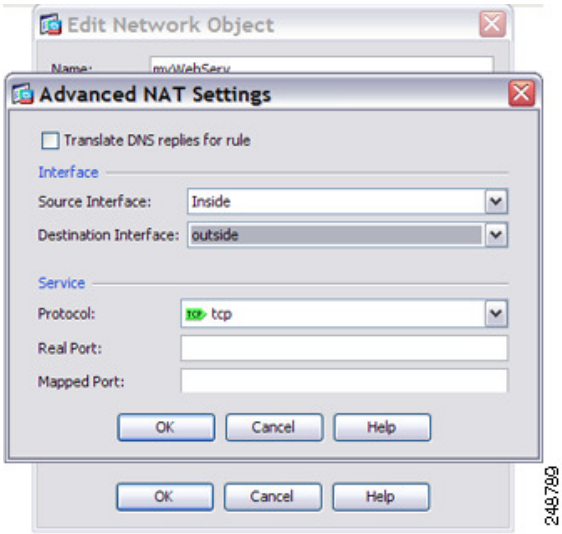
a) Choose **Add > Network Object**, name the new object, define the range of addresses, and click **OK**.



b) Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.



Step 5 Click **Advanced** and configure the real and mapped interfaces.

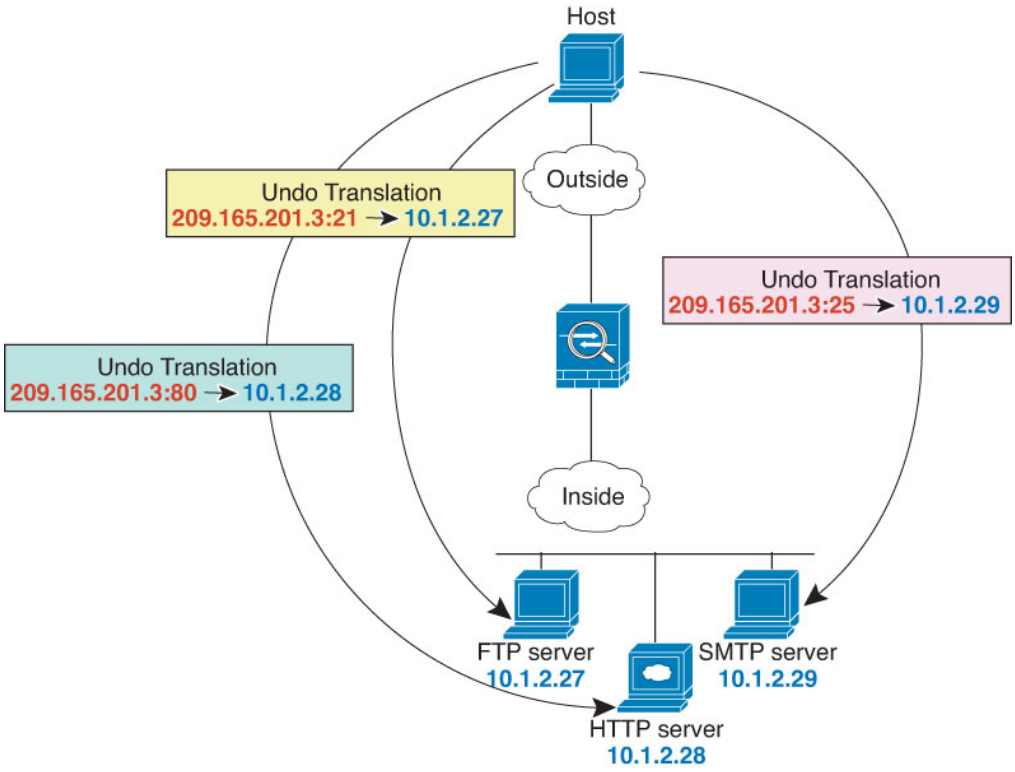


Step 6 Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

Single Address for FTP, HTTP, and SMTP (Static NAT-with-Port-Translation)

The following static NAT-with-port-translation example provides a single address for remote users to access FTP, HTTP, and SMTP. These servers are actually different devices on the real network, but for each server, you can specify static NAT-with-port-translation rules that use the same mapped IP address, but different ports.

Figure 31: Static NAT-with-Port-Translation



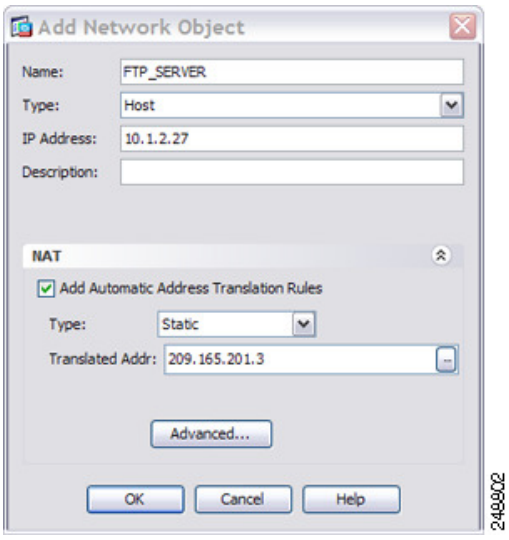
Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT**.

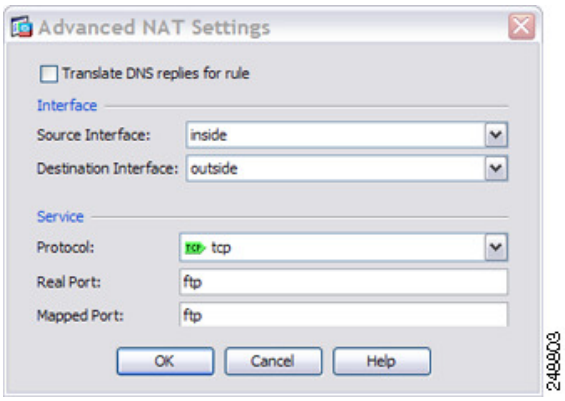
Step 2 Configure the static network object NAT with port translation rule for the FTP server.

a) Choose **Add > Network Object NAT Rule**.

b) Name the new network object, define the FTP server address, enable static NAT, and enter the translated address.



c) Click **Advanced** and configure the real and mapped interfaces and port translation for FTP, mapping the FTP port to itself.



d) Click **OK**, then **OK** again to save the rule and return to the NAT page.

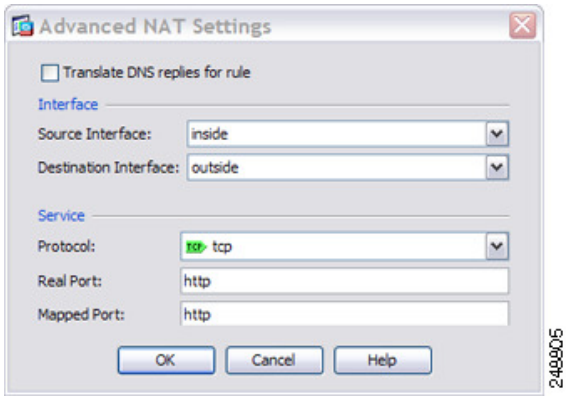
Step 3 Configure the static network object NAT with port translation rule for the HTTP server.

a) Choose **Add > Network Object NAT Rule**.

b) Name the new network object, define the HTTP server address, enable static NAT, and enter the translated address.



c) Click **Advanced** and configure the real and mapped interfaces and port translation for HTTP, mapping the HTTP port to itself.

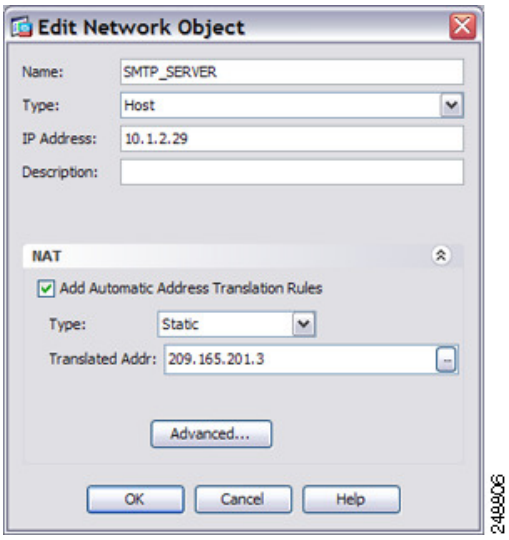


d) Click **OK**, then **OK** again to save the rule and return to the NAT page.

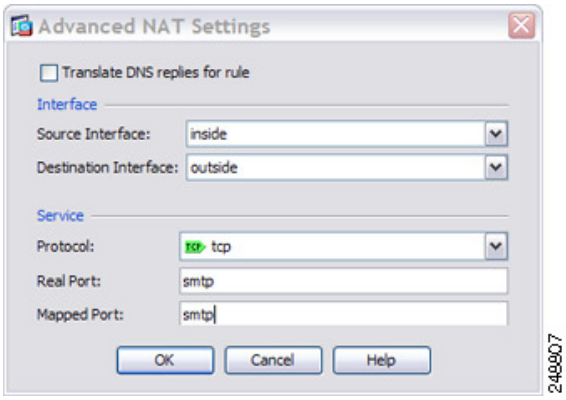
Step 4 Configure the static network object NAT with port translation rule for the SMTP server.

a) Choose **Add > Network Object NAT Rule**.

b) Name the new network object, define the SMTP server address, enable static NAT, and enter the translated address.



c) Click **Advanced** and configure the real and mapped interfaces and port translation for SMTP, mapping the SMTP port to itself.



d) Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

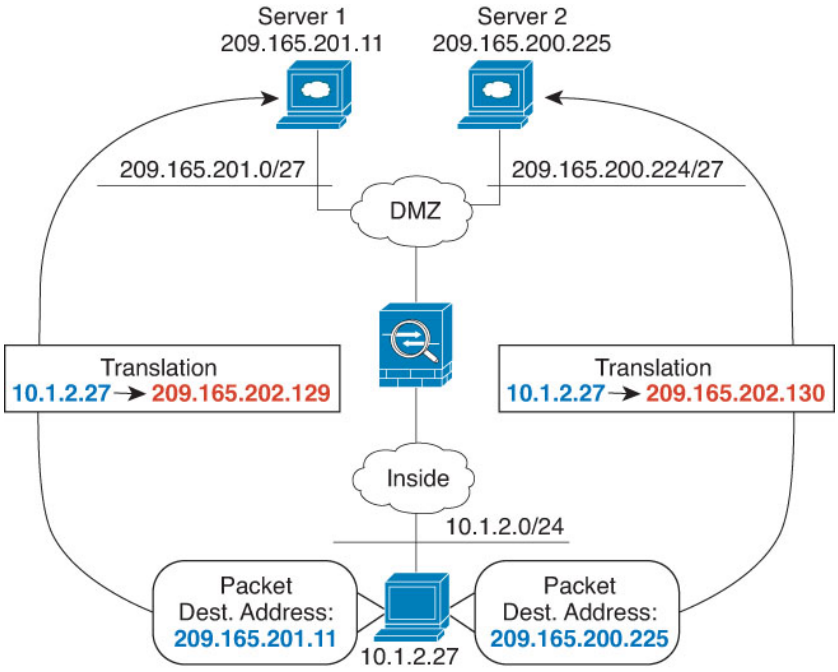
Examples for Twice NAT

This section includes the following configuration examples:

Different Translation Depending on the Destination (Dynamic Twice PAT)

The following figure shows a host on the 10.1.2.0/24 network accessing two different servers. When the host accesses the server at 209.165.201.11, the real address is translated to 209.165.202.129:*port*. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130:*port*.

Figure 32: Twice NAT with Different Destination Addresses

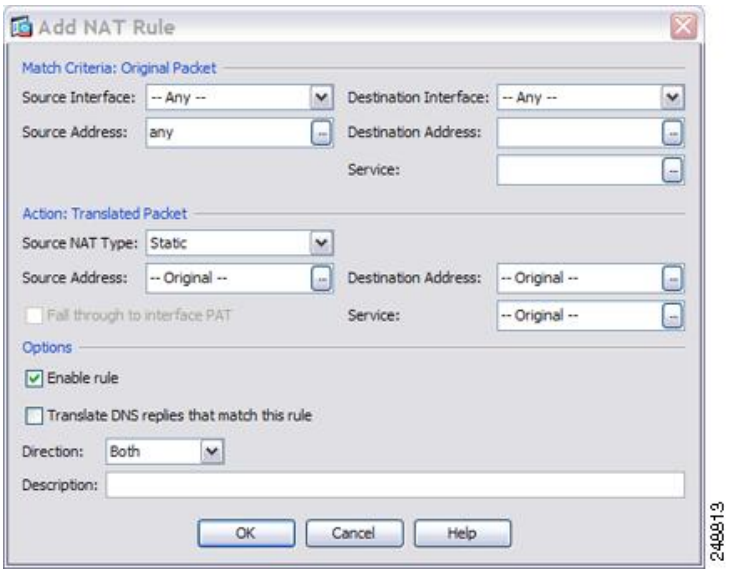


Procedure

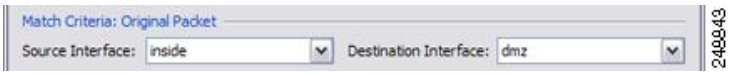
Step 1 On the **Configuration** > **Firewall** > **NAT Rules** page, click **Add** > **Add NAT Rule Before Network Object NAT Rules** to add a NAT rule for traffic from the inside network to DMZ network 1.

If you want to add a NAT rule to section 3, after the network object NAT rules, choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.



Step 2 Set the source and destination interfaces.



Step 3 For the Original Source Address, click the browse button to add a new network object for the inside network in the Browse Original Source Address dialog box.

a) Select **Add > Network Object**.

b) Define the inside network addresses, and click **OK**.



c) Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.



Step 4 For the Original Destination Address, click the browse button to add a new network object for DMZ network 1 in the Browse Original Destination Address dialog box.

a) Select **Add > Network Object**.

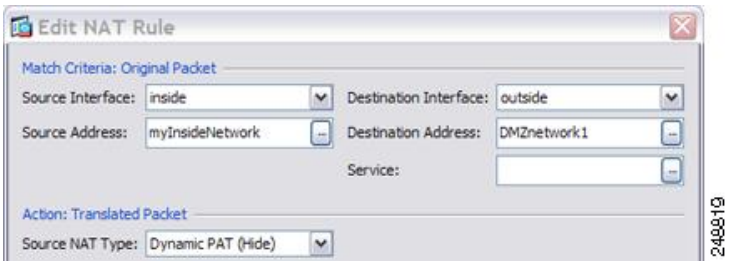
b) Define the DMZ network 1 addresses, and click **OK**.



c) Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.



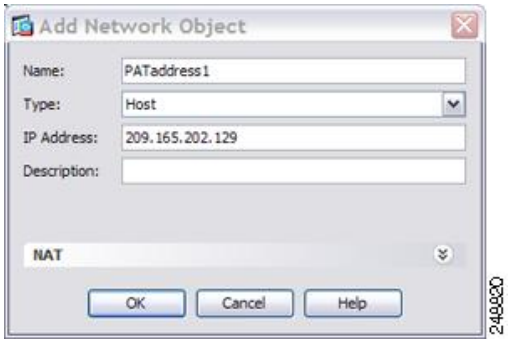
Step 5 Set the NAT Type to **Dynamic PAT (Hide)**:



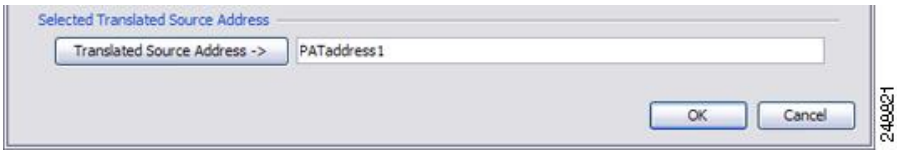
Step 6 For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

a) Select **Add > Network Object**.

b) Define the PAT address, and click **OK**.

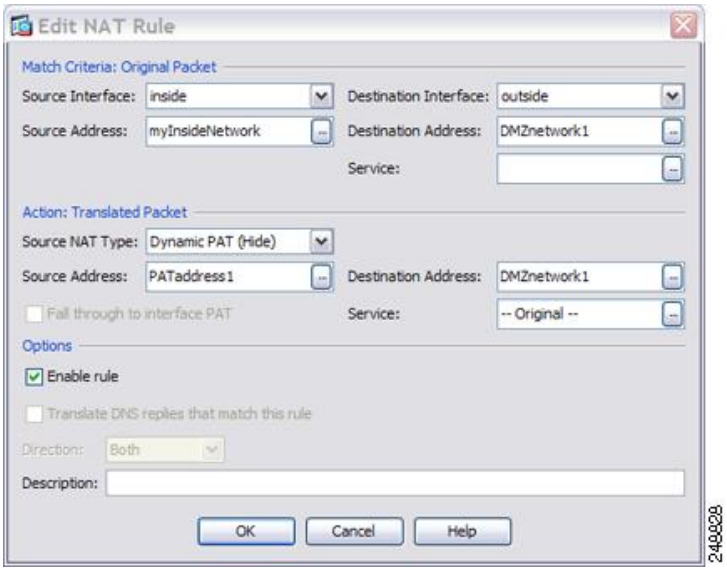


c) Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.



Step 7 For the Translated Destination Address, type the name of the Original Destination Address (DMZnetwork1) or click the browse button to choose it.

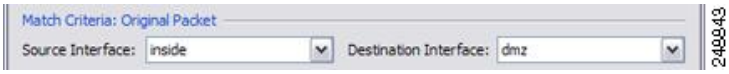
Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the Original and Translated destination addresses.



Step 8 Click **OK** to add the rule to the NAT table.

Step 9 Click **Add** > **Add NAT Rule Before Network Object NAT Rules** or **Add NAT Rule After Network Object NAT Rules** to add a NAT rule for traffic from the inside network to DMZ network 2.

Step 10 Set the source and destination interfaces.



Step 11 For the Original Source Address, type the name of the inside network object (myInsideNetwork) or click the browse button to choose it.

Step 12 For the Original Destination Address, click the browse button to add a new network object for DMZ network 2 in the Browse Original Destination Address dialog box.

a) Select **Add > Network Object**.

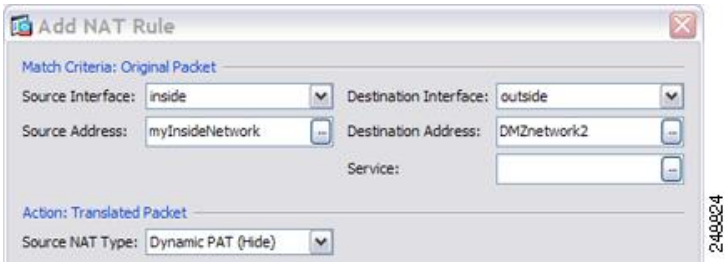
b) Define the DMZ network 2 addresses, and click **OK**.



c) Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.



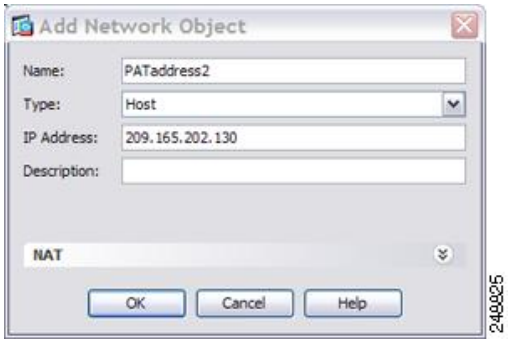
Step 13 Set the NAT Type to **Dynamic PAT (Hide)**:



Step 14 For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

a) Select **Add > Network Object**.

b) Define the PAT address, and click **OK**.

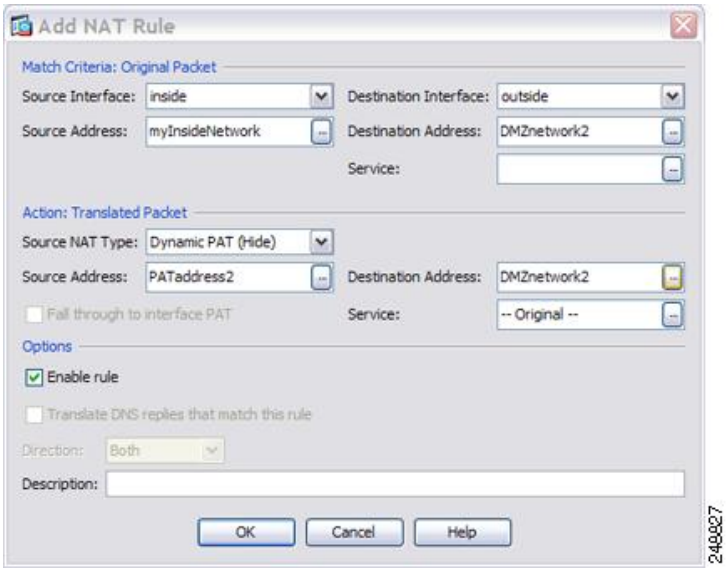


c) Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.



Step 15 For the Translated Destination Address, type the name of the Original Destination Address (DMZnetwork2) or click the browse button to choose it.

Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the Original and Translated destination addresses.



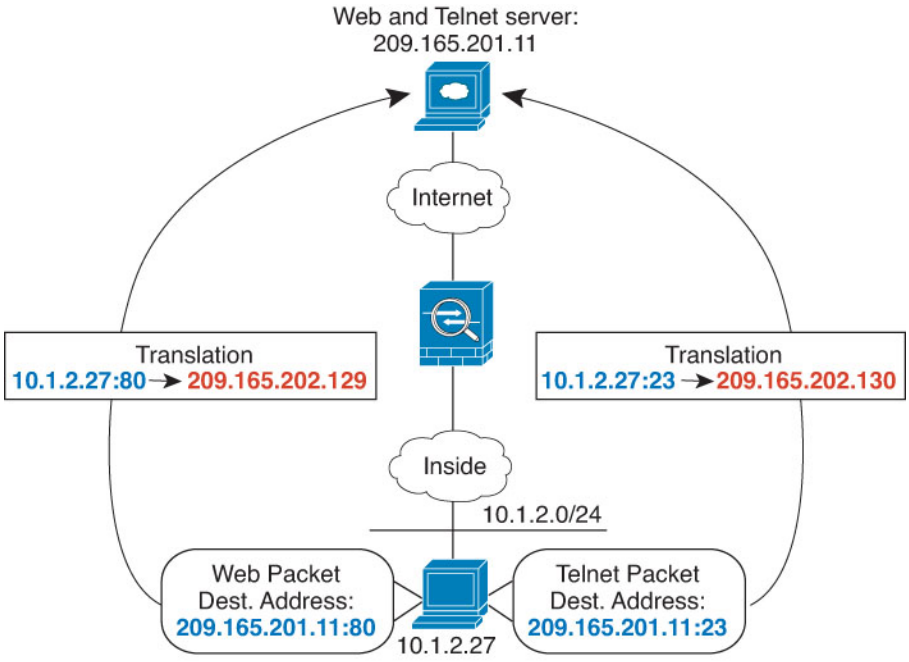
Step 16 Click **OK** to add the rule to the NAT table.

Step 17 Click **Apply**.

Different Translation Depending on the Destination Address and Port (Dynamic PAT)

The following figure shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for Telnet services, the real address is translated to 209.165.202.129:*port*. When the host accesses the same server for web services, the real address is translated to 209.165.202.130:*port*.

Figure 33: Twice NAT with Different Destination Ports

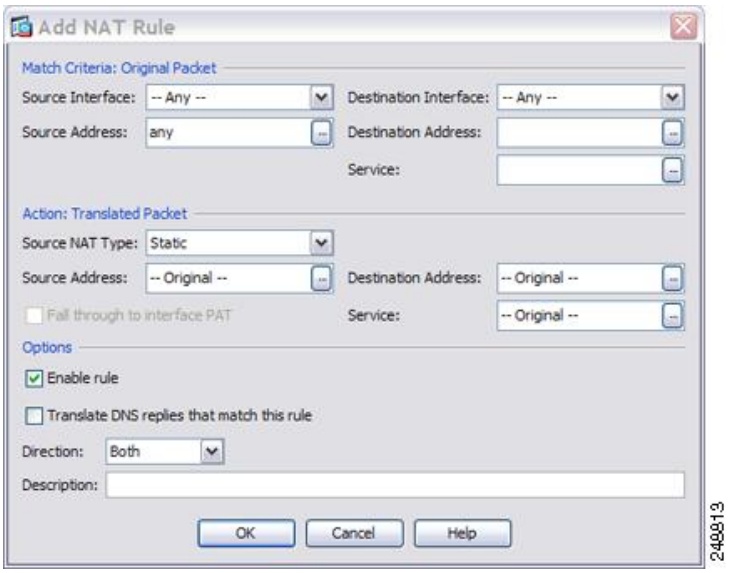


Procedure

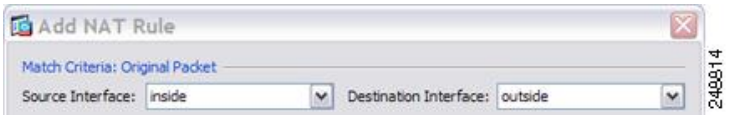
Step 1 On the **Configuration** > **Firewall** > **NAT Rules** page, click **Add** > **Add NAT Rule Before Network Object NAT Rules** to add a NAT rule for traffic from the inside network to the Telnet server.

If you want to add a NAT rule to section 3, after the network object NAT rules, choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.



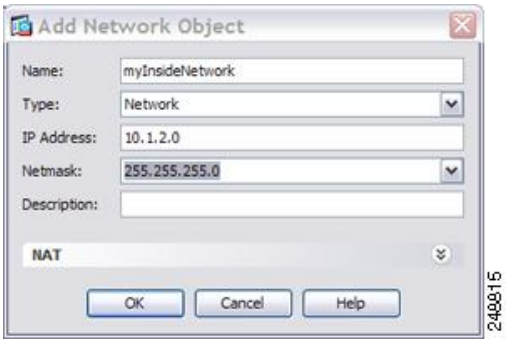
Step 2 Set the source and destination interfaces.



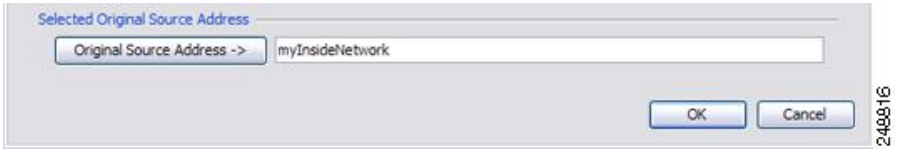
Step 3 For the Original Source Address, click the browse button to add a new network object for the inside network in the Browse Original Source Address dialog box.

a) Select **Add > Network Object**.

b) Define the inside network addresses, and click **OK**.

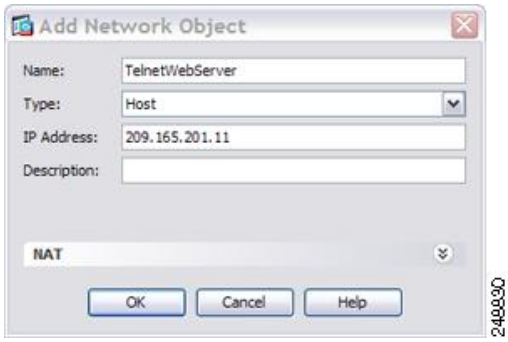


c) Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.



Step 4 For the Original Destination Address, click the browse button to add a new network object for the Telnet/Web server in the Browse Original Destination Address dialog box.

a) Select **Add > Network Object**.
b) Define the server address, and click **OK**.

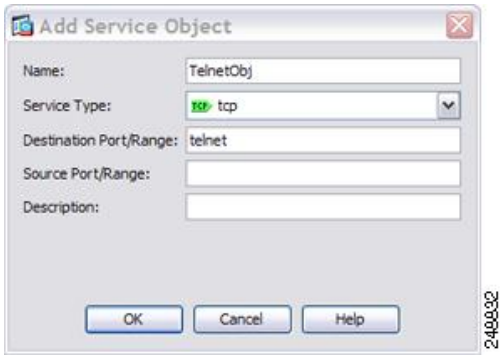


c) Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.



Step 5 For the Original Service, click the browse button to add a new service object for Telnet in the Browse Original Service dialog box.

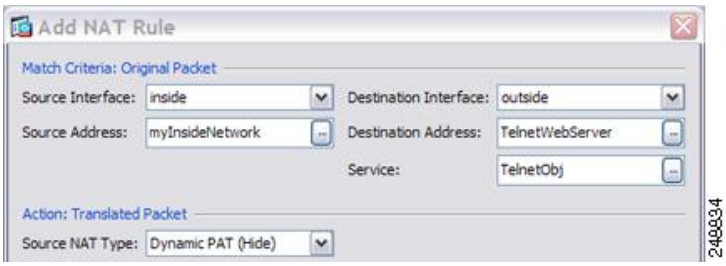
a) Select **Add > Service Object**.
b) Define the protocol and port, and click **OK**.



c) Choose the new service object by double-clicking it. Click **OK** to return to the NAT configuration.



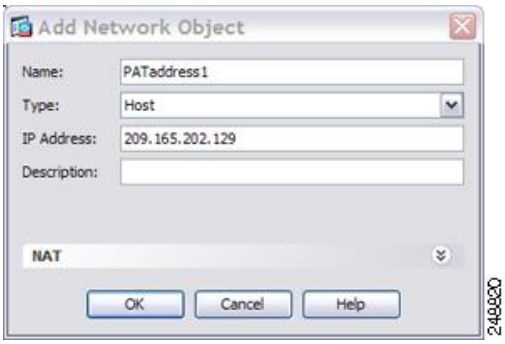
Step 6 Set the NAT Type to **Dynamic PAT (Hide)**:



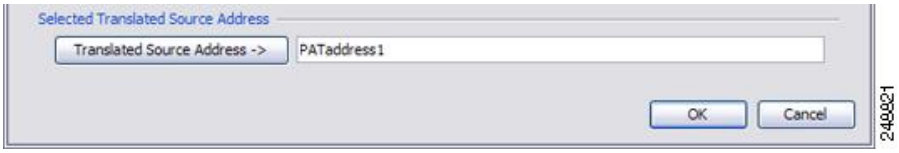
Step 7 For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

a) Select **Add > Network Object**.

b) Define the PAT address, and click **OK**.

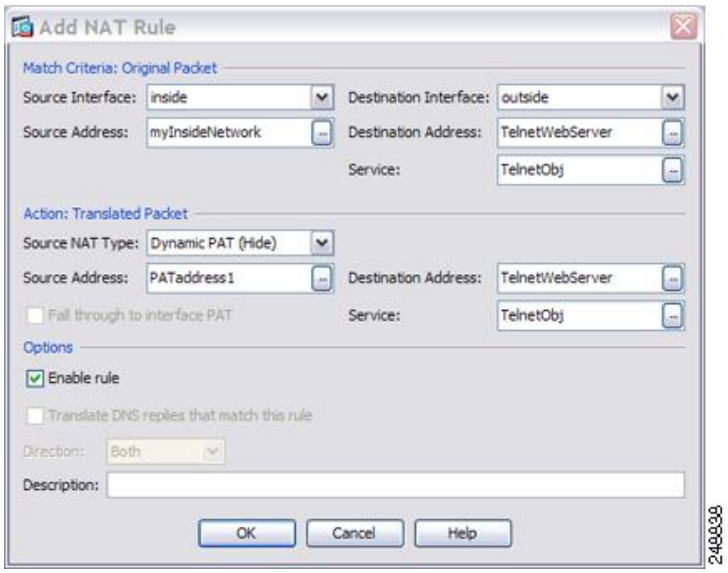


c) Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.



Step 8 For the Translated Destination Address, type the name of the Original Destination Address (TelnetWebServer) or click the browse button to choose it.

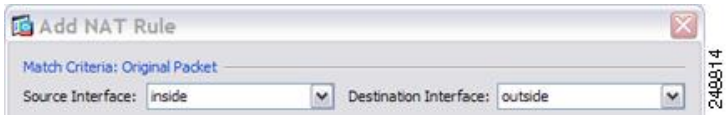
Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the Original and Translated destination addresses.



Step 9 Click **OK** to add the rule to the NAT table.

Step 10 Click **Add** > **Add NAT Rule Before Network Object NAT Rules** or **Add NAT Rule After Network Object NAT Rules** to add a NAT rule for traffic from the inside network to the web server.

Step 11 Set the real and mapped interfaces.



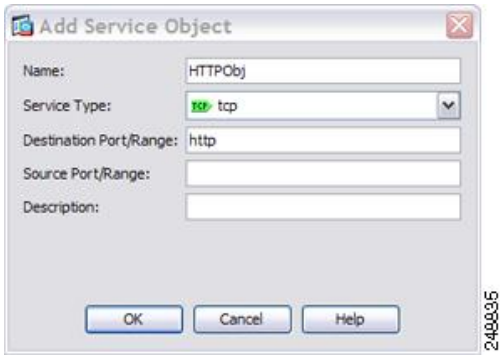
Step 12 For the Original Source Address, type the name of the inside network object (myInsideNetwork) or click the browse button to choose it.

Step 13 For the Original Destination Address, type the name of the Telnet/web server network object (TelnetWebServer) or click the browse button to choose it.

Step 14 For the Original Service, click the browse button to add a new service object for HTTP in the Browse Original Service dialog box.

a) Select **Add > Service Object**.

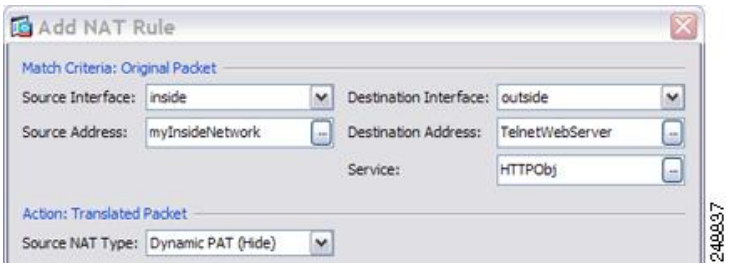
b) Define the protocol and port, and click **OK**.



c) Choose the new service object by double-clicking it. Click **OK** to return to the NAT configuration.



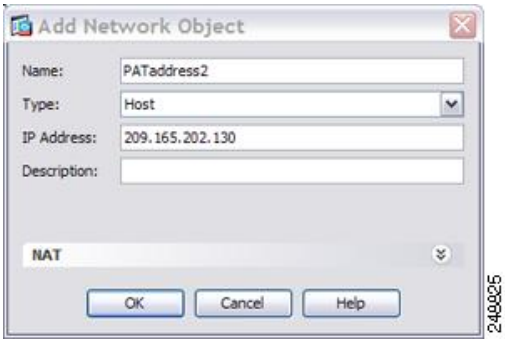
Step 15 Set the NAT Type to **Dynamic PAT (Hide)**:



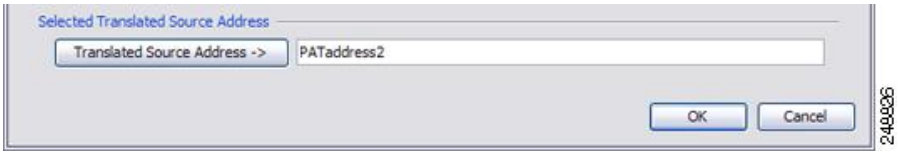
Step 16 For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

a) Select **Add > Network Object**.

b) Define the PAT address, and click **OK**.

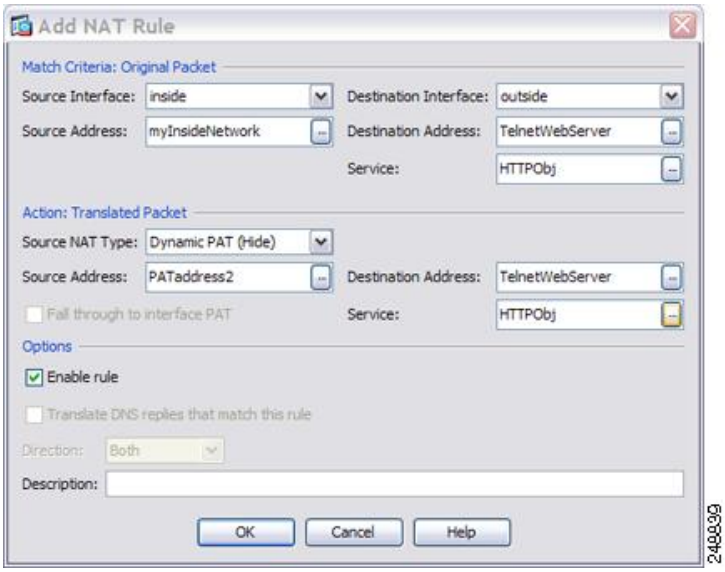


c) Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.



Step 17 For the Translated Destination Address, type the name of the Original Destination Address (TelnetWebServer) or click the browse button to choose it.

Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the Original and Translated destination addresses.



Step 18 Click **OK** to add the rule to the NAT table.

Step 19 Click **Apply**.

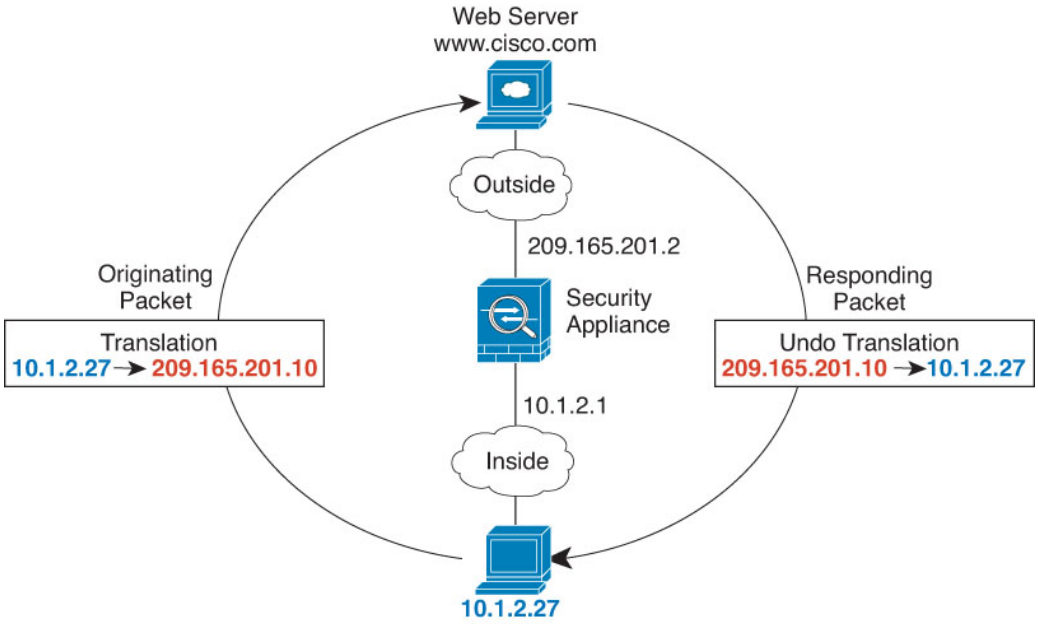
NAT in routed and transparent mode

You can configure NAT in both routed and transparent firewall mode. The following sections describe typical usage for each firewall mode.

NAT in routed mode

The following figure shows a typical NAT example in routed mode, with a private network on the inside.

Figure 34: NAT Example: Routed Mode



1. When the inside host at 10.1.2.27 sends a packet to a web server, the real source address of the packet, 10.1.2.27, is translated to a mapped address, 209.165.201.10.
2. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the ASA receives the packet because the ASA performs proxy ARP to claim the packet.
3. The ASA then changes the translation of the mapped address, 209.165.201.10, back to the real address, 10.1.2.27, before sending it to the host.

NAT in transparent mode or within a bridge group

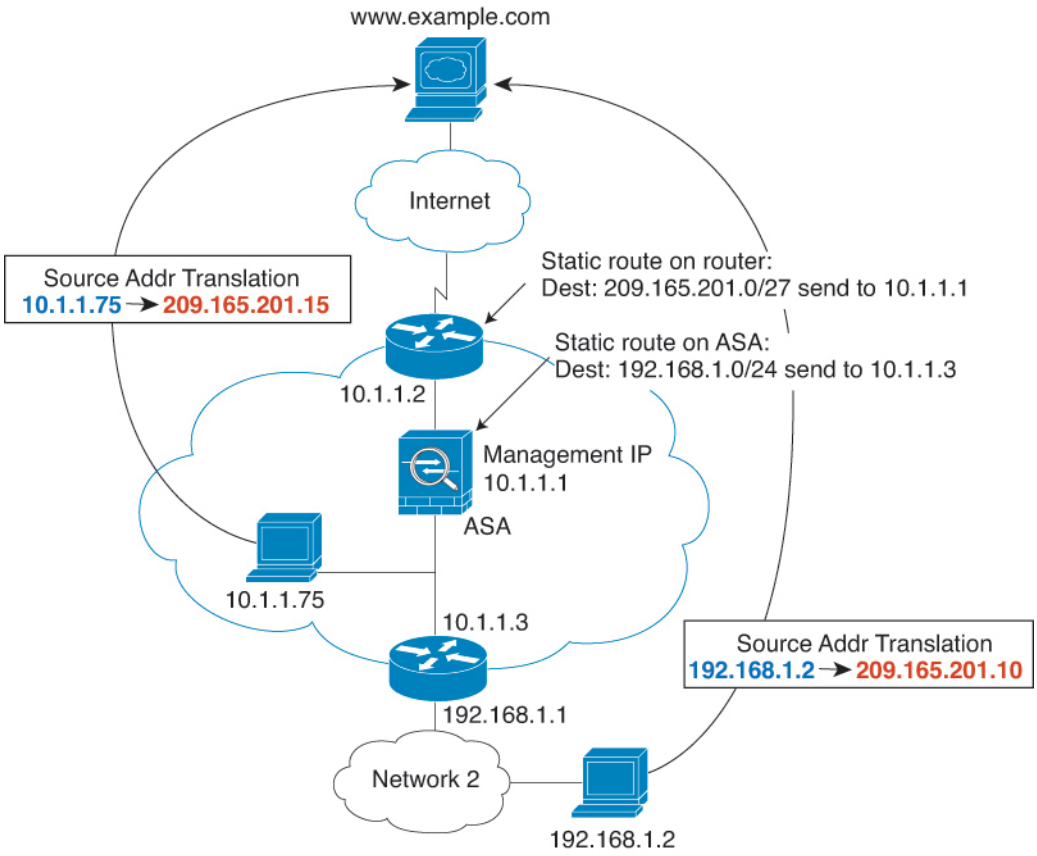
Using NAT in transparent mode eliminates the need for the upstream or downstream routers to perform NAT for their networks. It can perform a similar function within a bridge group in routed mode.

NAT in transparent mode, or in routed mode between members of the same bridge group, has the following requirements and limitations:

- You cannot configure interface PAT when the mapped address is a bridge group member interface, because there is no IP address attached to the interface.
- ARP inspection is not supported. Moreover, if for some reason a host on one side of the ASA sends an ARP request to a host on the other side of the ASA, and the initiating host real address is mapped to a different address on the same subnet, then the real address remains visible in the ARP request.
- Translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

The following figure shows a typical NAT scenario in transparent mode, with the same network on the inside and outside interfaces. The transparent firewall in this scenario is performing the NAT service so that the upstream router does not have to perform NAT.

Figure 35: NAT Example: Transparent Mode



1. When the inside host at 10.1.1.75 sends a packet to a web server, the real source address of the packet, 10.1.1.75, is changed to a mapped address, 209.165.201.15.
2. When the server responds, it sends the response to the mapped address, 209.165.201.15, and the ASA receives the packet because the upstream router includes this mapped network in a static route directed to the ASA management IP address.
3. The ASA then undoes the translation of the mapped address, 209.165.201.15, back to the real address, 10.1.1.1.75. Because the real address is directly-connected, the ASA sends it directly to the host.
4. For host 192.168.1.2, the same process occurs, except for returning traffic, the ASA looks up the route in its routing table and sends the packet to the downstream router at 10.1.1.3 based on the ASA static route for 192.168.1.0/24.

Routing NAT Packets

The ASA needs to be the destination for any packets sent to the mapped address. The ASA also needs to determine the egress interface for any packets it receives destined for mapped addresses. This section describes how the ASA handles accepting and delivering packets with NAT.

Mapped Addresses and Routing

When you translate the real address to a mapped address, the mapped address you choose determines how to configure routing, if necessary, for the mapped address.

See additional guidelines about mapped IP addresses in Additional guidelines for NAT, on page 156.

The following topics explain the mapped address types.

Addresses on the same network as the mapped interface

If you use addresses on the same network as the destination (mapped) interface, the ASA uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the ASA does not have to be the gateway for any additional networks. This solution is ideal if the outside network contains an adequate number of free addresses, a consideration if you are using a 1:1 translation like dynamic NAT or static NAT. Dynamic PAT greatly extends the number of translations you can use with a small number of addresses, so even if the available addresses on the outside network is small, this method can be used. For PAT, you can even use the IP address of the mapped interface.



Note If you configure the mapped interface to be any interface, and you specify a mapped address on the same network as one of the mapped interfaces, then if an ARP request for that mapped address comes in on a *different* interface, then you need to manually configure an ARP entry for that network on the ingress interface, specifying its MAC address. Typically, if you specify any interface for the mapped interface, then you use a unique network for the mapped addresses, so this situation would not occur. Select **Configuration** > **Device Management** > **Advanced** > **ARP** > **ARP Static Table** to configure ARP.

Addresses on a unique network

If you need more addresses than are available on the destination (mapped) interface network, you can identify addresses on a different subnet. The upstream router needs a static route for the mapped addresses that points to the ASA.

Alternatively for routed mode, you can configure a static route on the ASA for the mapped addresses using any IP address on the destination network as the gateway, and then redistribute the route using your routing protocol. For example, if you use NAT for the inside network (10.1.1.0/24) and use the mapped IP address 209.165.201.5, then you can configure a static route for 209.165.201.5 255.255.255.255 (host address) to the 10.1.1.99 gateway that can be redistributed.

```
route inside 209.165.201.5 255.255.255.255 10.1.1.99
```

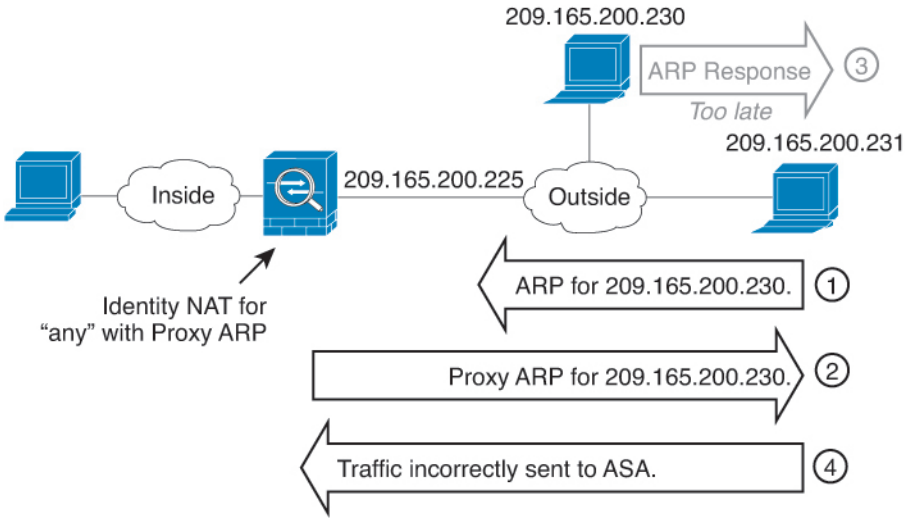
For transparent mode, if the real host is directly-connected, configure the static route on the upstream router to point to the ASA: specify the bridge group IP address. For remote hosts in transparent mode, in the static route on the upstream router, you can alternatively specify the downstream router IP address.

The same address as the real address (identity NAT)

The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. You can also disable proxy ARP for regular static NAT if desired, in which case you need to be sure to have proper routes on the upstream router.

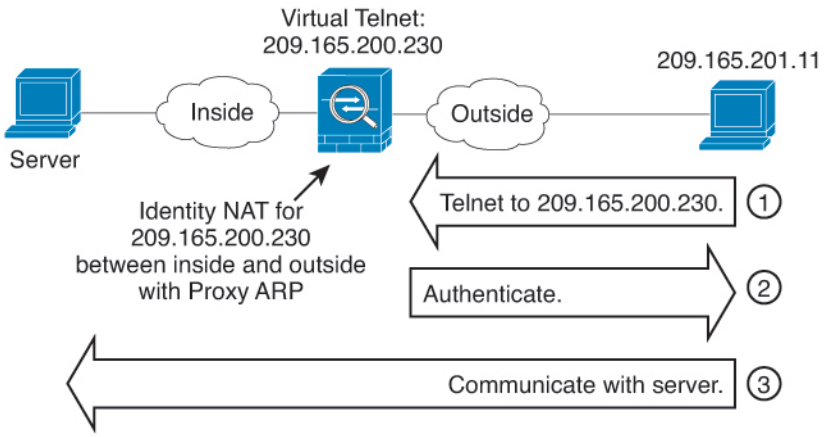
Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues. For example, if you configure a broad identity NAT rule for "any" IP address, then leaving proxy ARP enabled can cause problems for hosts on the network directly connected to the mapped interface. In this case, when a host on the mapped network wants to communicate with another host on the same network, then the address in the ARP request matches the NAT rule (which matches "any" address). The ASA will then proxy ARP for the address, even though the packet is not actually destined for the ASA. (Note that this problem occurs even if you have a twice NAT rule; although the NAT rule must match both the source and destination addresses, the proxy ARP decision is made only on the "source" address). If the ASA ARP response is received before the actual host ARP response, then traffic will be mistakenly sent to the ASA.

Figure 36: Proxy ARP Problems with Identity NAT



In rare cases, you need proxy ARP for identity NAT; for example for virtual Telnet. When using AAA for network access, a host needs to authenticate with the ASA using a service like Telnet before any other traffic can pass. You can configure a virtual Telnet server on the ASA to provide the necessary login. When accessing the virtual Telnet address from the outside, you must configure an identity NAT rule for the address specifically for the proxy ARP functionality. Due to internal processes for virtual Telnet, proxy ARP lets the ASA keep traffic destined for the virtual Telnet address rather than send the traffic out the source interface according to the NAT rule. (See the following figure).

Figure 37: Proxy ARP and Virtual Telnet



Transparent Mode Routing Requirements for Remote Networks

When you use NAT in transparent mode, some types of traffic require static routes. See the general operations configuration guide for more information.

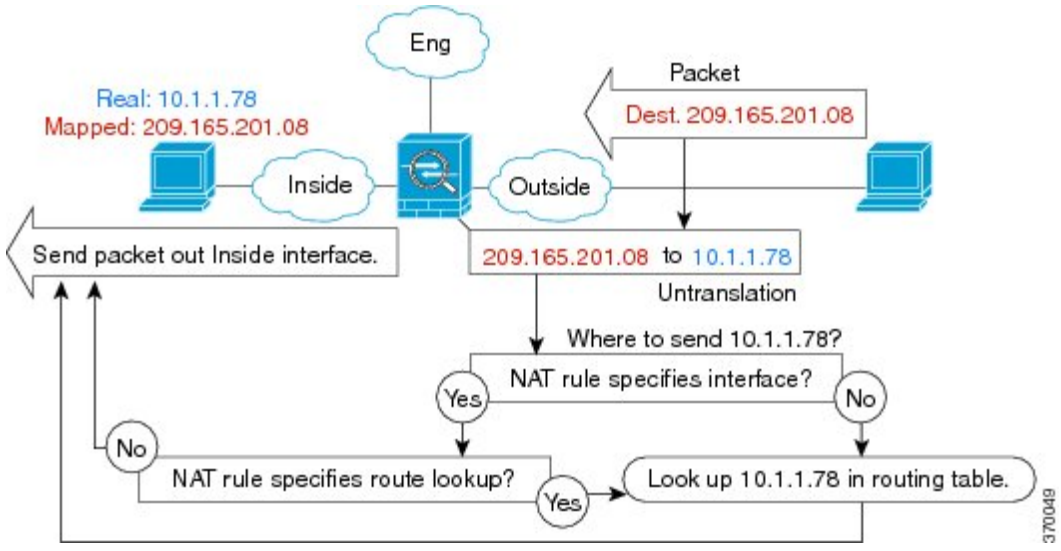
Determining the Egress Interface

When you use NAT and the ASA receives traffic for a mapped address, then the ASA untranslates the destination address according to the NAT rule, and then it sends the packet on to the real address. The ASA determines the egress interface for the packet in the following ways:

- Bridge group interfaces in Transparent mode or Routed Mode—The ASA determines the egress interface for the real address by using the NAT rule; you must specify the source and destination bridge group member interfaces as part of the NAT rule.
- Regular interfaces in Routed mode—The ASA determines the egress interface in one of the following ways:
  - You configure the interface in the NAT rule—The ASA uses the NAT rule to determine the egress interface. (8.3(1) through 8.4(1)) The only exception is for identity NAT, which always uses a route lookup, regardless of the NAT configuration. (8.4(2) and later) For identity NAT, the default behavior is to use the NAT configuration. However, you have the option to always use a route lookup instead. In certain scenarios, a route lookup override is required.
  - You do not configure the interface in the NAT rule—The ASA uses a route lookup to determine the egress interface.

The following figure shows the egress interface selection method in routed mode. In almost all cases, a route lookup is equivalent to the NAT rule interface, but in some configurations, the two methods might differ.

Figure 38: Routed Mode Egress Interface Selection with NAT



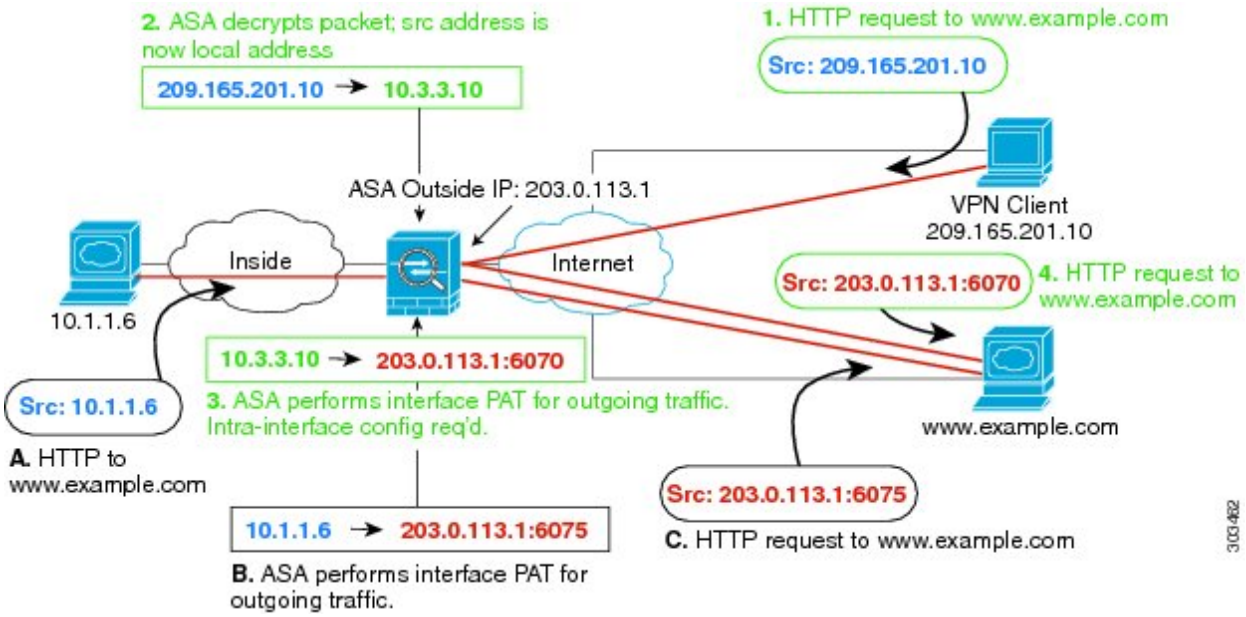
NAT for VPN

The following topics explain NAT usage with the various types of VPN.

NAT and Remote Access VPN

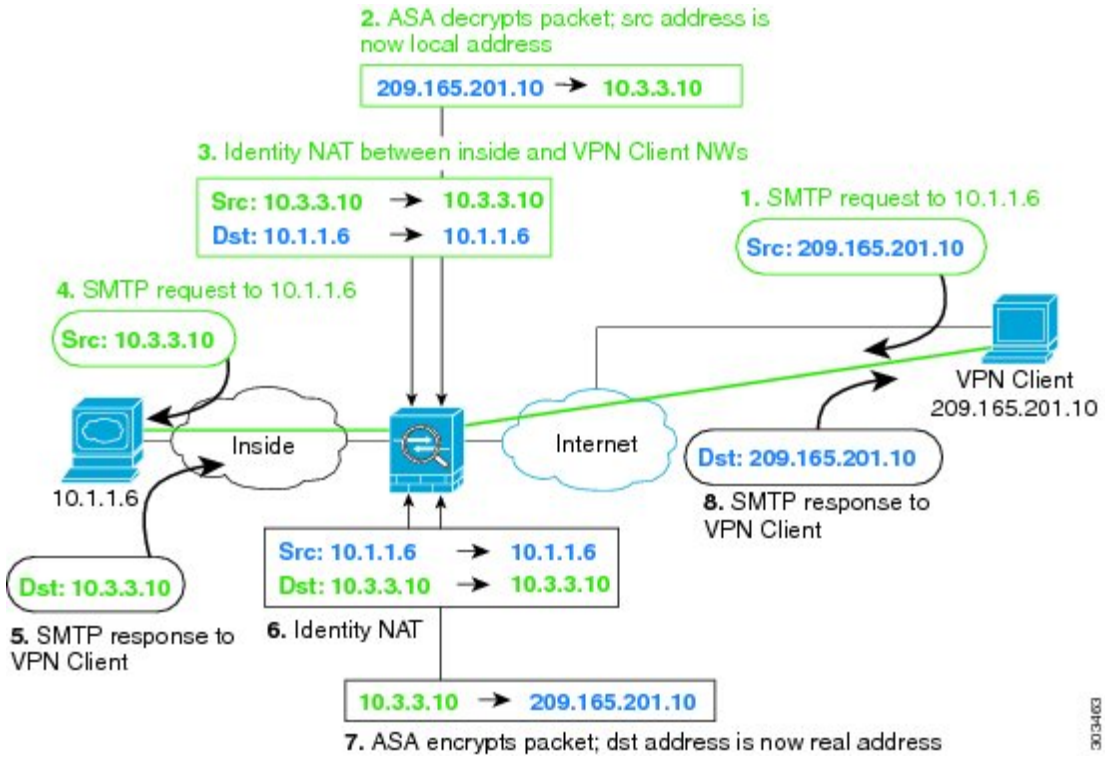
The following figure shows both an inside server (10.1.1.6) and a VPN client (209.165.201.10) accessing the Internet. Unless you configure split tunneling for the VPN client (where only specified traffic goes through the VPN tunnel), then Internet-bound VPN traffic must also go through the ASA. When the VPN traffic enters the ASA, the ASA decrypts the packet; the resulting packet includes the VPN client local address (10.3.3.10) as the source. For both inside and VPN client local networks, you need a public IP address provided by NAT to access the Internet. The below example uses interface PAT rules. To allow the VPN traffic to exit the same interface it entered, you also need to enable intra-interface communication (also known as "hairpin" networking).

Figure 39: Interface PAT for Internet-Bound VPN Traffic (Intra-Interface)



The following figure shows a VPN client that wants to access an inside mail server. Because the ASA expects traffic between the inside network and any outside network to match the interface PAT rule you set up for Internet access, traffic from the VPN client (10.3.3.10) to the SMTP server (10.1.1.6) will be dropped due to a reverse path failure: traffic from 10.3.3.10 to 10.1.1.6 does not match a NAT rule, but returning traffic from 10.1.1.6 to 10.3.3.10 *should* match the interface PAT rule for outgoing traffic. Because forward and reverse flows do not match, the ASA drops the packet when it is received. To avoid this failure, you need to exempt the inside-to-VPN client traffic from the interface PAT rule by using an identity NAT rule between those networks. Identity NAT simply translates an address to the same address.

Figure 40: Identity NAT for VPN Clients



See the following sample NAT configuration for the above network:

```
! Enable hairpin for non-split-tunneled VPN client traffic:
same-security-traffic permit intra-interface

! Identify local VPN network, & perform object interface PAT when going to Internet:
object network vpn_local
subnet 10.3.3.0 255.255.255.0
nat (outside,outside) dynamic interface

! Identify inside network, & perform object interface PAT when going to Internet:
object network inside_nw
subnet 10.1.1.0 255.255.255.0
nat (inside,outside) dynamic interface

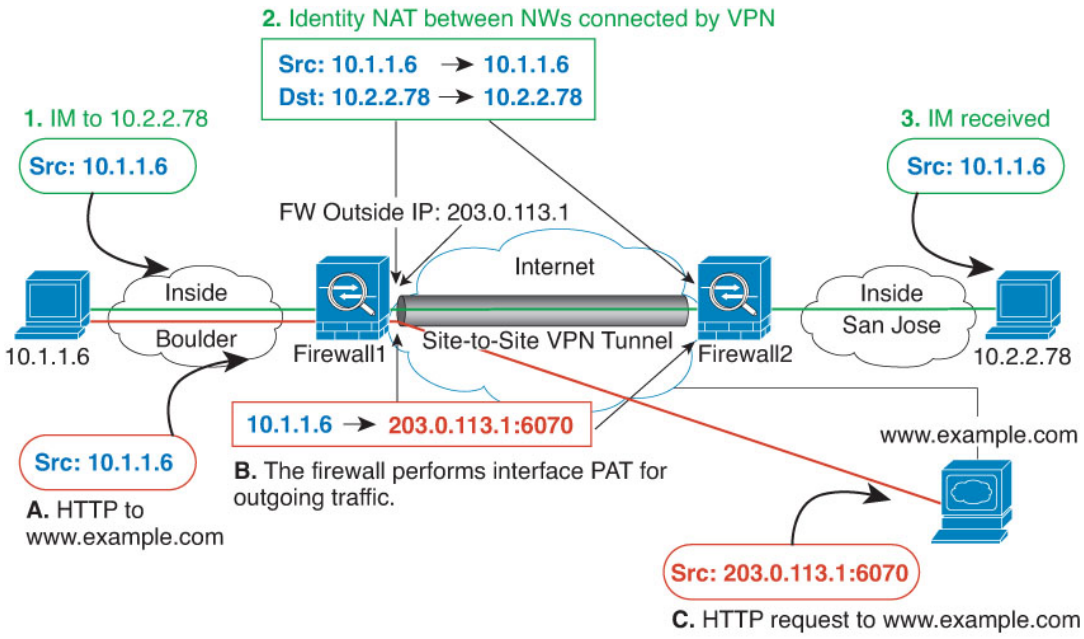
! Use twice NAT to pass traffic between the inside network and the VPN client without
! address translation (identity NAT):
nat (inside,outside) source static inside_nw inside_nw destination static vpn_local vpn_local
```

NAT and Site-to-Site VPN

The following figure shows a site-to-site tunnel connecting the Boulder and San Jose offices. For traffic that you want to go to the Internet (for example from 10.1.1.6 in Boulder to www.example.com), you need a public IP address provided by NAT to access the Internet. The below example uses interface PAT rules. However, for traffic that you want to go over the VPN tunnel (for example from 10.1.1.6 in Boulder to 10.2.2.78 in San

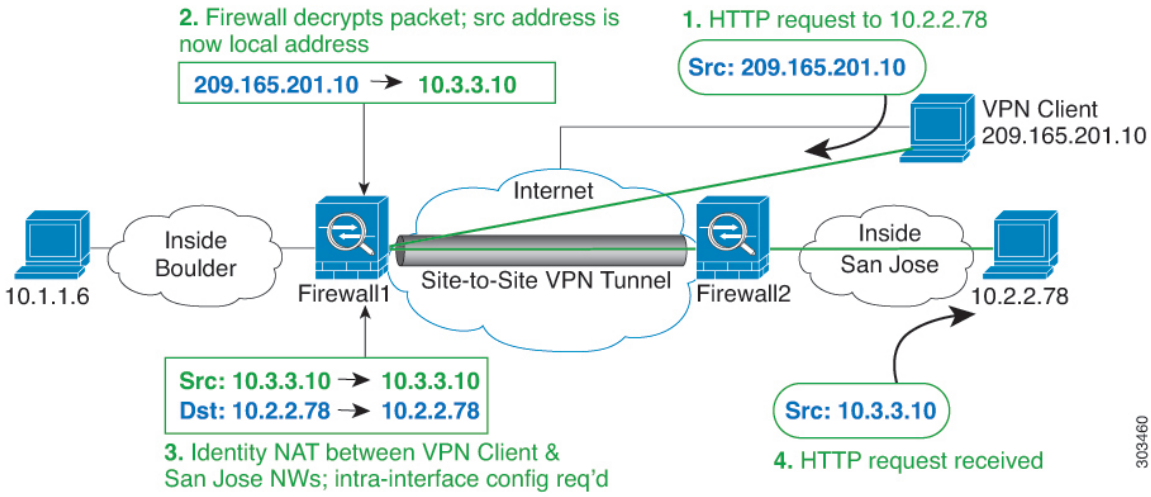
Jose), you do not want to perform NAT; you need to exempt that traffic by creating an identity NAT rule. Identity NAT simply translates an address to the same address.

Figure 41: Interface PAT and Identity NAT for Site-to-Site VPN



The following figure shows a VPN client connected to Firewall1 (Boulder), with a Telnet request for a server (10.2.2.78) accessible over a site-to-site tunnel between Firewall1 and Firewall2 (San Jose). Because this is a hairpin connection, you need to enable intra-interface communication, which is also required for non-split-tunneled Internet-bound traffic from the VPN client. You also need to configure identity NAT between the VPN client and the Boulder & San Jose networks, just as you would between any networks connected by VPN to exempt this traffic from outbound NAT rules.

Figure 42: VPN Client Access to Site-to-Site VPN



See the following sample NAT configuration for Firewall1 (Boulder) for the second example:

```
! Enable hairpin for VPN client traffic:
same-security-traffic permit intra-interface

! Identify local VPN network, & perform object interface PAT when going to Internet:
object network vpn_local
subnet 10.3.3.0 255.255.255.0
nat (outside,outside) dynamic interface

! Identify inside Boulder network, & perform object interface PAT when going to Internet:
object network boulder_inside
subnet 10.1.1.0 255.255.255.0
nat (inside,outside) dynamic interface

! Identify inside San Jose network for use in twice NAT rule:
object network sanjose_inside
subnet 10.2.2.0 255.255.255.0

! Use twice NAT to pass traffic between the Boulder network and the VPN client without
! address translation (identity NAT):
nat (inside,outside) source static boulder_inside boulder_inside
destination static vpn_local vpn_local

! Use twice NAT to pass traffic between the Boulder network and San Jose without
! address translation (identity NAT):
nat (inside,outside) source static boulder_inside boulder_inside
destination static sanjose_inside sanjose_inside

! Use twice NAT to pass traffic between the VPN client and San Jose without
! address translation (identity NAT):
nat (outside,outside) source static vpn_local vpn_local
destination static sanjose_inside sanjose_inside
```

See the following sample NAT configuration for Firewall2 (San Jose):

```
! Identify inside San Jose network, & perform object interface PAT when going to Internet:
object network sanjose_inside
subnet 10.2.2.0 255.255.255.0
nat (inside,outside) dynamic interface

! Identify inside Boulder network for use in twice NAT rule:
object network boulder_inside
subnet 10.1.1.0 255.255.255.0

! Identify local VPN network for use in twice NAT rule:
object network vpn_local
subnet 10.3.3.0 255.255.255.0

! Use twice NAT to pass traffic between the San Jose network and Boulder without
! address translation (identity NAT):
nat (inside,outside) source static sanjose_inside sanjose_inside
destination static boulder_inside boulder_inside

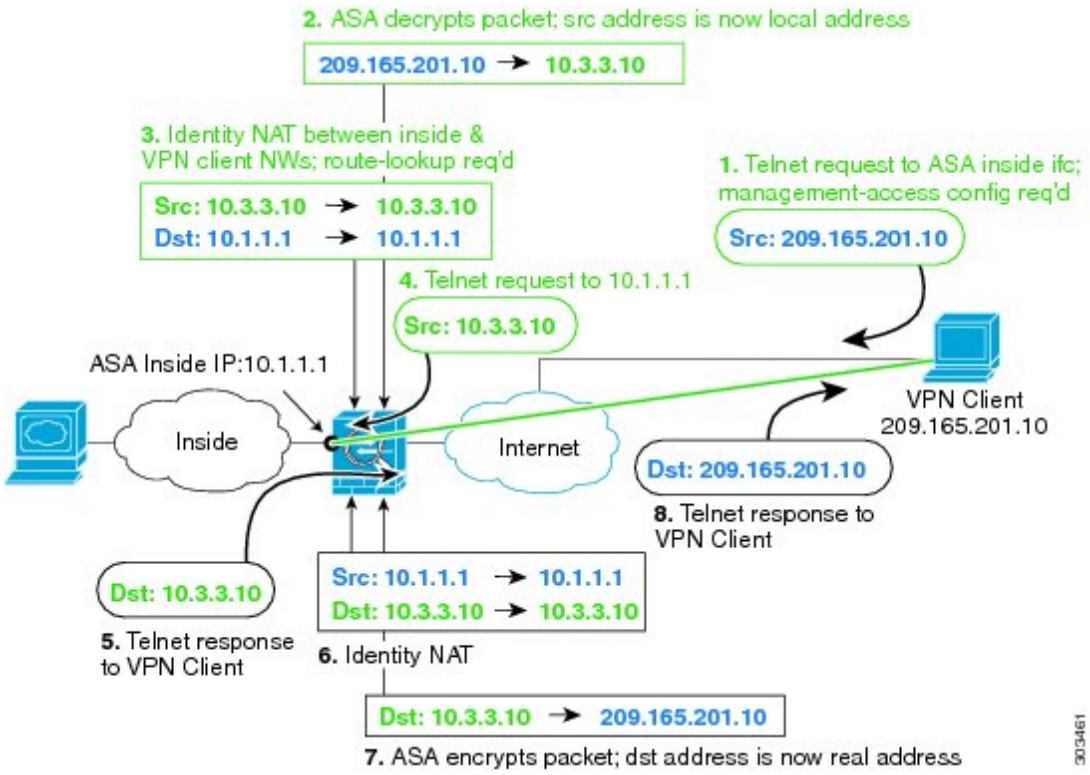
! Use twice NAT to pass traffic between the San Jose network and the VPN client without
! address translation (identity NAT):
nat (inside,outside) source static sanjose_inside sanjose_inside
destination static vpn_local vpn_local
```

NAT and VPN Management Access

When using VPN, you can allow management access to an interface other than the one from which you entered the ASA. For example, if you enter the ASA from the outside interface, the management-access feature lets you connect to the inside interface using ASDM, SSH, Telnet, or SNMP; or you can ping the inside interface.

The following figure shows a VPN client Telnetting to the ASA inside interface. When you use a management-access interface, and you configure identity NAT according to NAT and Remote Access VPN, on page 247 or NAT and Site-to-Site VPN, on page 248, you must configure NAT with the route lookup option. Without route lookup, the ASA sends traffic out the interface specified in the NAT command, regardless of what the routing table says; in the below example, the egress interface is the inside interface. You do not want the ASA to send the management traffic out to the inside network; it will never return to the inside interface IP address. The route lookup option lets the ASA send the traffic directly to the inside interface IP address instead of to the inside network. For traffic from the VPN client to a host on the inside network, the route lookup option will still result in the correct egress interface (inside), so normal traffic flow is not affected. See the Determining the Egress Interface, on page 246 for more information about the route lookup option.

Figure 43: VPN Management Access



See the following sample NAT configuration for the above network:

```
! Enable hairpin for non-split-tunneled VPN client traffic:
same-security-traffic permit intra-interface

! Enable management access on inside ifc:
management-access inside

! Identify local VPN network, & perform object interface PAT when going to Internet:
object network vpn_local
```

```
subnet 10.3.3.0 255.255.255.0
nat (outside,outside) dynamic interface

! Identify inside network, & perform object interface PAT when going to Internet:
object network inside_nw
subnet 10.1.1.0 255.255.255.0
nat (inside,outside) dynamic interface

! Use twice NAT to pass traffic between the inside network and the VPN client without
! address translation (identity NAT), w/route-lookup:
nat (outside,inside) source static vpn_local vpn_local
destination static inside_nw inside_nw route-lookup
```

Troubleshooting NAT and VPN

See the following monitoring tools for troubleshooting NAT issues with VPN:

- Packet tracer—When used correctly, a packet tracer shows which NAT rules a packet is hitting.
- **show nat detail**—Shows hit counts and untranslated traffic for a given NAT rule.
- **show conn all**—Lets you see active connections including to and from the box traffic.

To familiarize yourself with a non-working configuration vs. a working configuration, you can perform the following steps:

1. Configure VPN without identity NAT.
2. Enter **show nat detail** and **show conn all**.
3. Add the identity NAT configuration.
4. Repeat **show nat detail** and **show conn all**.

Translating IPv6 networks

In cases where you need to pass traffic between IPv6-only and IPv4-only networks, you need to use NAT to convert between the address types. Even with two IPv6 networks, you might want to hide internal addresses from the outside network.

You can use the following translation types with IPv6 networks:

- NAT64, NAT46—Translates IPv6 packets into IPv4 and vice versa. You need to define two policies, one for the IPv6 to IPv4 translation, and one for the IPv4 to IPv6 translation. Although you can accomplish this with a single twice NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a twice NAT rule when you specify a destination, creating two network object NAT rules is the better solution.



Note NAT46 supports static mappings only.

- NAT66—Translates IPv6 packets to a different IPv6 address. We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT.

Note NAT64 and NAT 46 are possible on standard routed interfaces only. NAT66 is possible on both routed and bridge group member interfaces.

NAT64/46: translating IPv6 addresses to IPv4

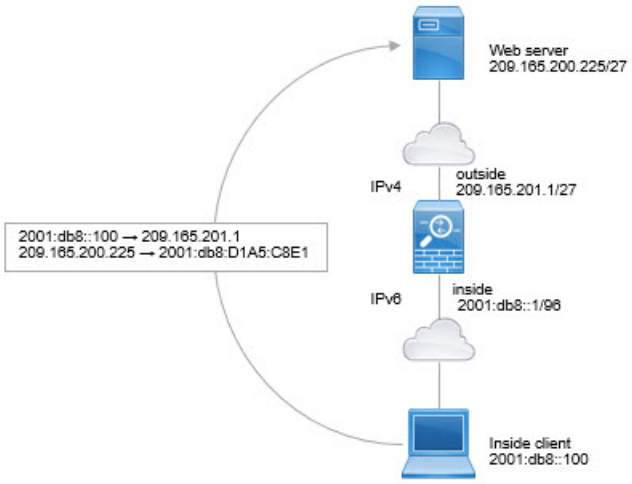
When traffic goes from an IPv6 network to an IPv4-only network, you need to convert the IPv6 address to IPv4, and return traffic from IPv4 to IPv6. You need to define two address pools, an IPv4 address pool to bind IPv6 addresses in the IPv4 network, and an IPv6 address pool to bind IPv4 addresses in the IPv6 network.

- The IPv4 address pool for the NAT64 rule is normally small and typically might not have enough addresses to map one-to-one with the IPv6 client addresses. Dynamic PAT might more easily meet the possible large number of IPv6 client addresses compared to dynamic or static NAT.
- The IPv6 address pool for the NAT46 rule can be equal to or larger than the number of IPv4 addresses to be mapped. This allows each IPv4 address to be mapped to a different IPv6 address. NAT46 supports static mappings only, so you cannot use dynamic PAT.

You need to define two policies, one for the source IPv6 network, and one for the destination IPv4 network. Although you can accomplish this with a single twice NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a twice NAT rule when you specify a destination, creating two network object NAT rules is the better solution.

NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet

Following is a straight-forward example where you have an inside IPv6-only network, and you want to convert to IPv4 for traffic sent to the Internet. This example assumes you do not need DNS translation, so you can perform both the NAT64 and NAT46 translations in a single twice NAT rule.



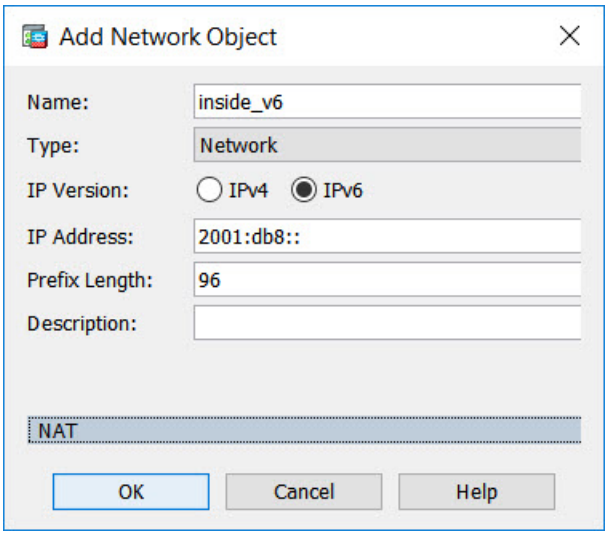
In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network.

Procedure

Step 1 Create a network object for the inside IPv6 network.

a) Choose **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups**.
b) Click **Add** > **Network Object**.
c) Configure an object with the following properties:

- **Name**—For example, inside_v6.
- **Type**—Select **Network**.
- **IP Version**—Select IPv6.
- **IP Address**—Enter 2001:db8::
- **Prefix Length**—Enter 96.



d) Click **OK**.

Step 2 Create the twice NAT rule to translate the IPv6 network to IPv4 and back again.

a) Choose **Configuration** > **Firewall** > **NAT Rules**.
b) Click **Add** > **Add NAT Rule Before "Network Object" NAT Rules**.
c) Configure the following **Match Criteria: Original Packet** options:

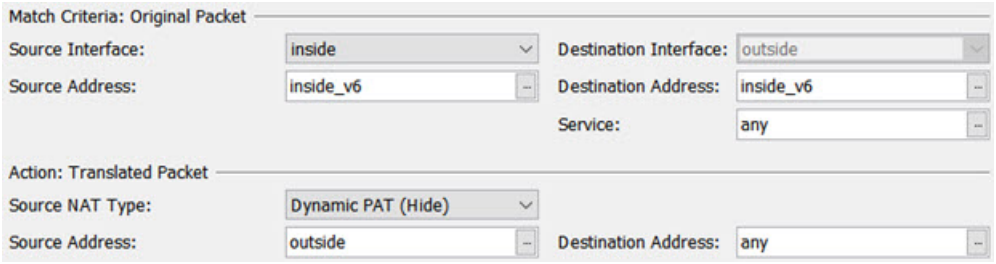
- **Source Interface**—Select inside.
- **Destination Interface**—Select outside.
- **Source Address**—Select the inside_v6 network object.
- **Destination Address**—Select the inside_v6 network object.
- **Service**—Keep the default, any.

d) Configure the following **Match Criteria: Translated Packet** options:

- **Source NAT Type**—Select Dynamic PAT (Hide).
- **Source Address**—Select the outside interface.
- **Destination Address**—Select any.

Leave the rest of the options at their default values.

The dialog box should look like the following:

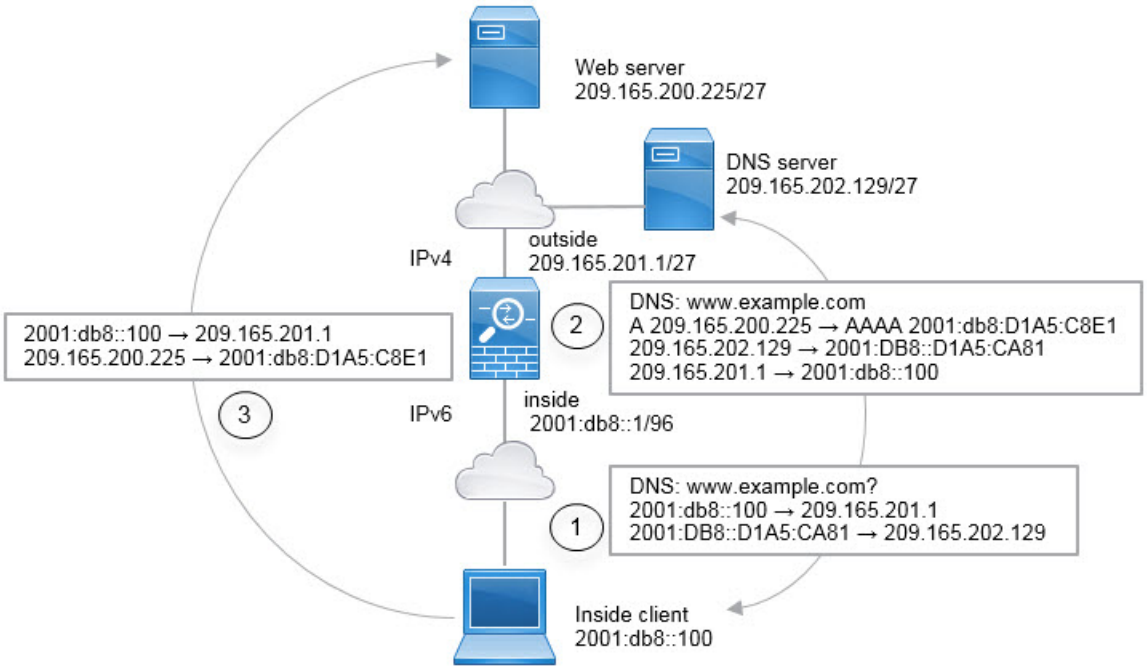


e) Click **OK**.

With this rule, any traffic from the 2001:db8::/96 subnet on the inside interface going to the outside interface gets a NAT64 PAT translation using the IPv4 address of the outside interface. Conversely, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method.

NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation

Following is a typical example where you have an inside IPv6-only network, but there are some IPv4-only services on the outside Internet that internal users need.



In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network. You enable DNS rewrite on the NAT46 rule, so that replies from the external DNS server can be converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.

Following is a typical sequence for a web request where a client at 2001:DB8::100 on the internal IPv6 network tries to open www.example.com.

1. The client's computer sends a DNS request to the DNS server at 2001:DB8::D1A5:CA81. The NAT rules make the following translations to the source and destination in the DNS request:
   - 2001:DB8::100 to a unique port on 209.165.201.1 (The NAT64 interface PAT rule.)
   - 2001:DB8::D1A5:CA81 to 209.165.202.129 (The NAT46 rule. D1A5:CA81 is the IPv6 equivalent of 209.165.202.129.)
2. The DNS server responds with an A record indicating that www.example.com is at 209.165.200.225. The NAT46 rule, with DNS rewrite enabled, converts the A record to the IPv6-equivalent AAAA record, and translates 209.165.200.225 to 2001:db8:D1A5:C8E1in the AAAA record. In addition, the source and destination addresses in the DNS response are untranslated:
   - 209.165.202.129 to 2001:DB8::D1A5:CA81
   - 209.165.201.1 to 2001:db8::100
3. The IPv6 client now has the IP address of the web server, and makes an HTTP request to www.example.com at 2001:db8:D1A5:C8E1. (D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225.) The source and destination of the HTTP request are translated:
   - 2001:DB8::100 to a unique port on 209.156.101.54 (The NAT64 interface PAT rule.)
   - 2001:db8:D1A5:C8E1 to 209.165.200.225 (The NAT46 rule.)

The following procedure explains how to configure this example.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**.

Step 2 Configure the NAT64 dynamic PAT rule for the inside IPv6 network.

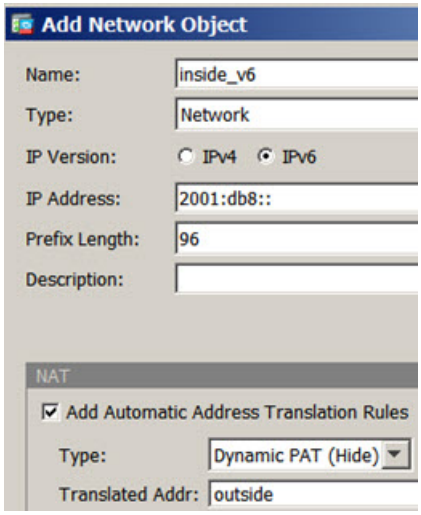
a) Choose **Add** > **Network Object NAT Rule**.

b) Configure the basic object properties:

- **Name**—For example, **inside_v6**.
- **Type**—Select **Network**.
- **IP Version**—Select **IPv6**.
- **IP Address**—Enter **2001:db8::**.
- **Prefix Length**—Enter **96**.

c) Select **Dynamic** or **Dynamic PAT (Hide)** for NAT Type.

d) For the **Translated Address**, click the browse button and select the "outside" interface.



e) Click the **Advanced** button and configure the following options:

- **Source Interface**—Select "inside."
- **Destination Interface**—The "outside" interface should already be selected.

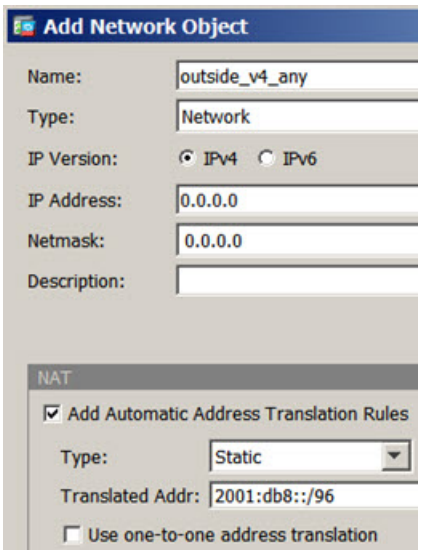
f) Click **OK** to save the advanced settings.

g) Click **OK** to add the NAT rule.

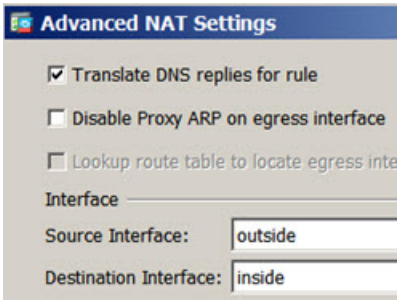
With this rule, any traffic from the 2001:db8::/96 subnet on the inside interface going to the outside interface gets a NAT64 PAT translation using the IPv4 address of the outside interface.

Step 3 Configure the static NAT46 rule for the outside IPv4 network.

a) Choose **Add** > **Network Object NAT Rule**.
b) Configure the basic object properties:
   - **Name**—For example, **outside_v4_any**.
   - **Type**—Select **Network**.
   - **IP Version**—Select **IPv4**.
   - **IP Address**—Enter **0.0.0.0**.
   - **Netmask**—Enter **0.0.0.0**.
c) Configure the NAT properties:
   - **NAT Type**—Select **Static**.
   - **Translated Address**—Enter **2001:db8::/96**.



d) Click the **Advanced** button and configure the following options:
   - **Translate DNS Replies for Rule**—Select this option.
   - **Source Interface**—Select “outside.”
   - **Destination Interface**—Select “inside.”



e) Click **OK** to save the advanced settings.
f) Click **OK** to add the NAT rule.

With this rule, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method. In addition, DNS responses are converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.

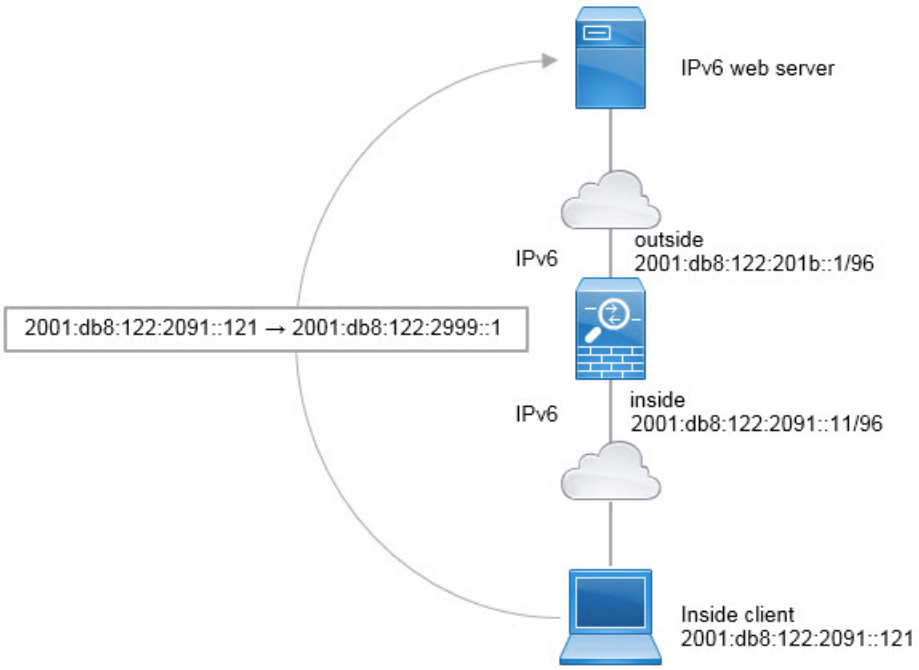
NAT66: translating IPv6 addresses to different IPv6 addresses

When going from an IPv6 network to another IPv6 network, you can translate the addresses to different IPv6 addresses on the outside network. We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT.

Because you are not translating between different address types, you need a single rule for NAT66 translations. You can easily model these rules using network object NAT. However, if you do not want to allow returning traffic, you can make the static NAT rule unidirectional using twice NAT only.

NAT66 Example, Static Translation between Networks

You can configure a static translation between IPv6 address pools using network object NAT. The following example explains how to convert inside addresses on the 2001:db8:122:2091::/96 network to outside addresses on the 2001:db8:122:2999::/96 network.



Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**.

Step 2 Configure the static NAT rule for the inside IPv6 network.

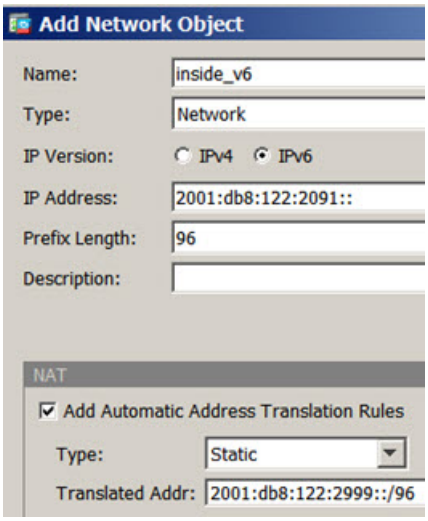
a) Choose **Add** > **Network Object NAT Rule**.

b) Configure the basic object properties:

- **Name**—For example, **inside_v6**.
- **Type**—Select **Network**.
- **IP Version**—Select **IPv6**.
- **IP Address**—Enter **2001:db8:122:2091::**.
- **Prefix Length**—Enter **96**.

c) Select **Static** for NAT Type.

d) For the **Translated Address**, enter **2001:db8:122:2999::/96**.



e) Click the **Advanced** button and configure the following options:

- **Source Interface**—Select “inside.”
- **Destination Interface**—Select “outside.”

f) Click **OK** to save the advanced settings.

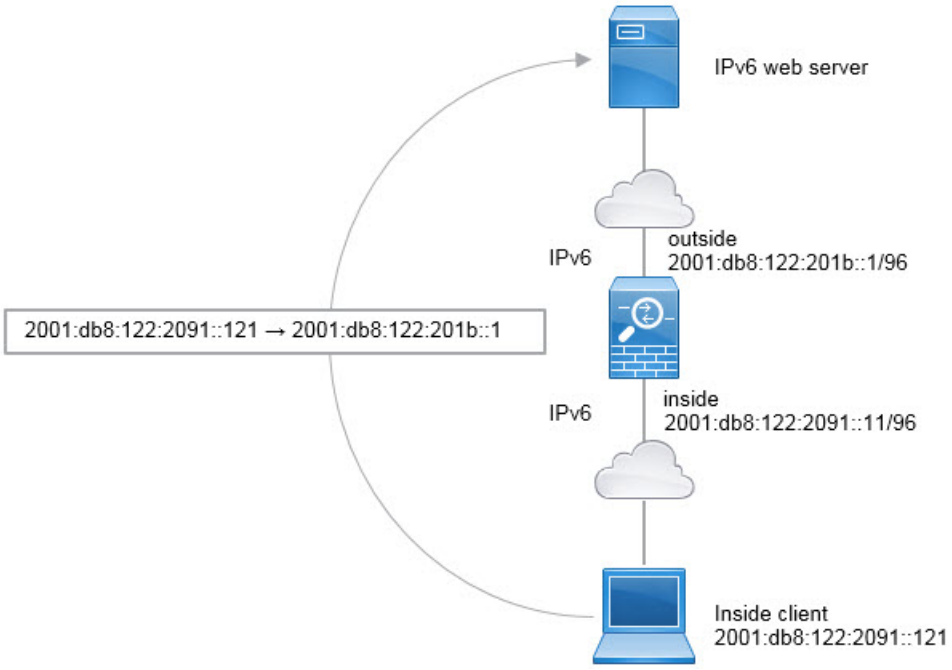
g) Click **OK** to add the NAT rule.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a static NAT66 translation to an address on the 2001:db8:122:2999::/96 network.

NAT66 Example, Simple IPv6 Interface PAT

A simple approach for implementing NAT66 is to dynamically assign internal addresses to different ports on the outside interface IPv6 address.

When you configure an interface PAT rule for NAT66, all the global addresses that are configured on that interface are used for PAT mapping. Link-local or site-local addresses for the interface are not used for PAT.



Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT Rules**.

Step 2 Configure the dynamic PAT rule for the inside IPv6 network.

a) Choose **Add** > **Network Object NAT Rule**.

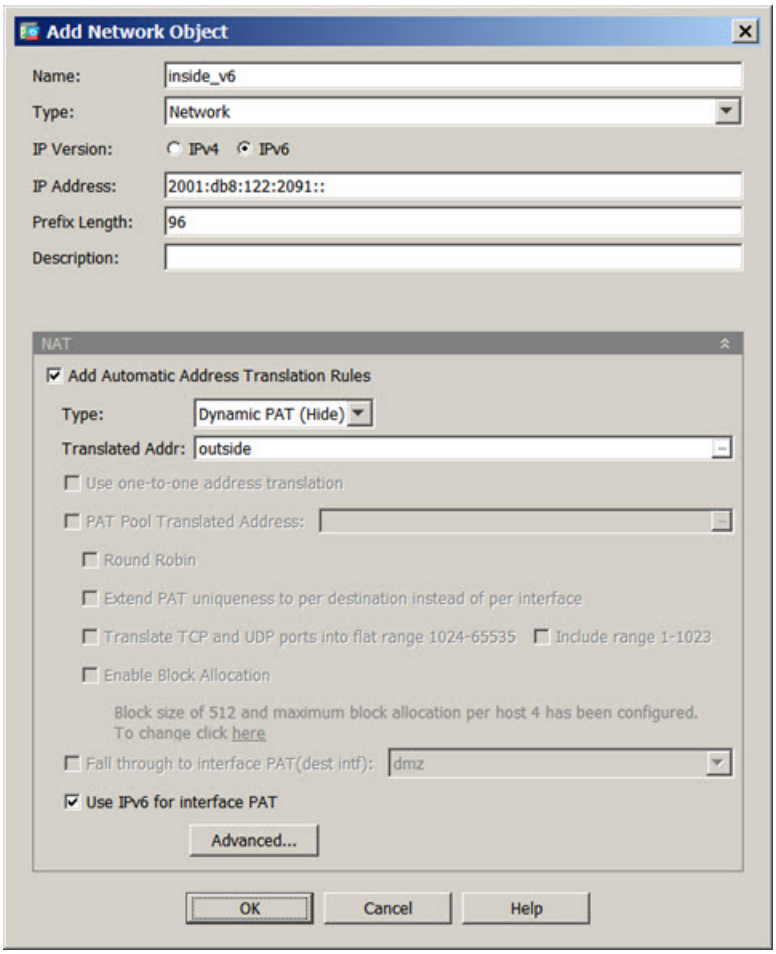
b) Configure the basic object properties:

- **Name**—For example, **inside_v6**.
- **Type**—Select **Network**.
- **IP Version**—Select **IPv6**.
- **IP Address**—Enter **2001:db8:122:2091::**.
- **Prefix Length**—Enter **96**.

c) Select **Dynamic** or **Dynamic PAT (Hide)** for NAT Type.

d) For the **Translated Address**, click the browse button and select the “outside” interface.

e) Select the **Use IPv6 for Interface PAT** option.



f) Click the **Advanced** button and configure the following options:

- **Source Interface**—Select "inside."
- **Destination Interface**—The "outside" interface should already be selected.

g) Click **OK** to save the advanced settings.

h) Click **OK** to add the NAT rule.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a NAT66 PAT translation to one of the IPv6 global addresses configured for the outside interface.

Rewriting DNS queries and responses using NAT

You might need to configure the ASA to modify DNS replies by replacing the address in the reply with an address that matches the NAT configuration. You can configure DNS modification when you configure each translation rule. DNS modification is also known as DNS doctoring.

This feature rewrites the address in DNS queries and replies that match a NAT rule (for example, the A record for IPv4, the AAAA record for IPv6, or the PTR record for reverse DNS queries). For DNS replies traversing from a mapped interface to any other interface, the record is rewritten from the mapped value to the real value. Inversely, for DNS replies traversing from any interface to a mapped interface, the record is rewritten from the real value to the mapped value. This feature works with NAT44,NAT 66, NAT46, and NAT64.

Following are the main circumstances when you would need to configure DNS rewrite on a NAT rule.

- The rule is NAT64 or NAT46, and the DNS server is on the outside network. You need DNS rewrite to convert between DNS A records (for IPv4) and AAAA records (for IPv6).
- The DNS server is on the outside, clients are on the inside, and some of the fully-qualified domain names that the clients use resolve to other inside hosts.
- The DNS server is on the inside and responds with private IP addresses, clients are on the outside, and the clients access fully-qualified domain names that point to servers that are hosted on the inside.

DNS rewrite limitations

Following are some limitations with DNS rewrite:

- DNS rewrite is not applicable for PAT because multiple PAT rules are applicable for each A or AAAA record, and the PAT rule to use is ambiguous.
- If you configure a twice NAT rule, you cannot configure DNS modification if you specify the destination address as well as the source address. These kinds of rules can potentially have a different translation for a single address when going to A vs. B. Therefore, the can not accurately match the IP address inside the DNS reply to the correct twice NAT rule; the DNS reply does not contain information about which source/destination address combination was in the packet that prompted the DNS request.
- You must enable DNS application inspection with DNS NAT rewrite enabled for NAT rules to rewrite DNS queries and responses. By default, DNS inspection with DNS NAT rewrite enabled is globally applied, so you probably do not need to change the inspection configuration.
- DNS rewrite is actually done on the xlate entry, not the NAT rule. Thus, if there is no xlate for a dynamic rule, rewrite cannot be done correctly. The same problem does not occur for static NAT.
- DNS rewrite does not rewrite DNS Dynamic Update messages (opcode 5).

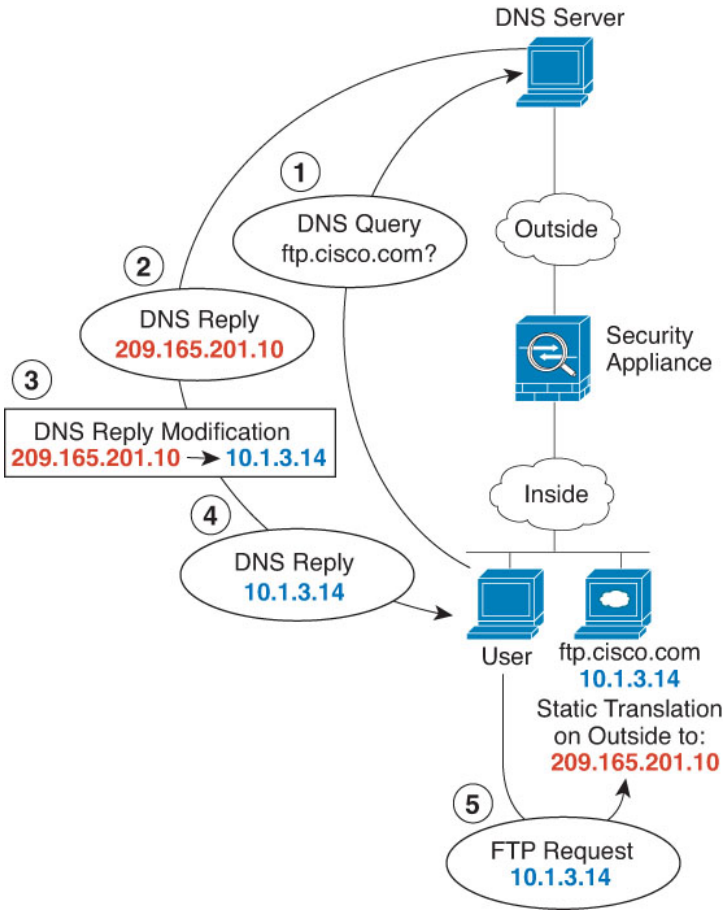
The following topics provide examples of DNS rewrite in NAT rules.

DNS Reply Modification, DNS Server on Outside

The following figure shows a DNS server that is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure NAT to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network.

In this case, you want to enable DNS reply modification on this static rule so that inside users who have access to ftp.cisco.com using the real address receive the real address from the DNS server, and not the mapped address.

When an inside host sends a DNS request for the address of ftp.cisco.com, the DNS server replies with the mapped address (209.165.201.10). The system refers to the static rule for the inside server and translates the address inside the DNS reply to 10.1.3.14. If you do not enable DNS reply modification, then the inside host attempts to send traffic to 209.165.201.10 instead of accessing ftp.cisco.com directly.

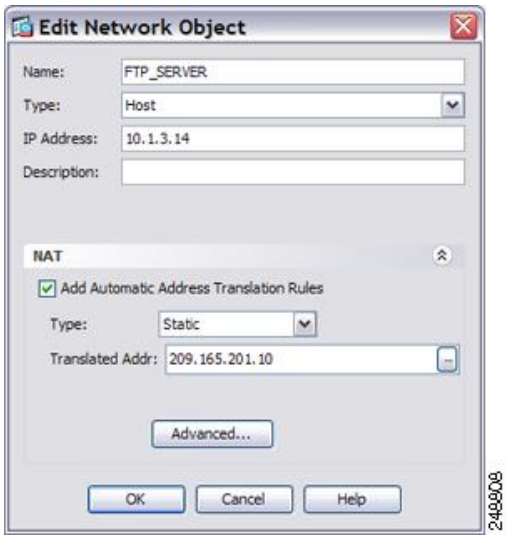


Procedure

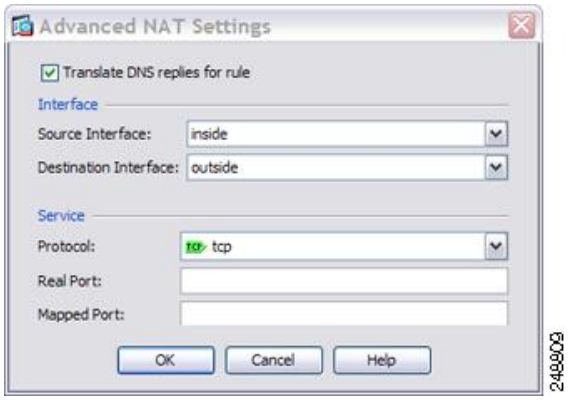
Step 1 Choose **Configuration** > **Firewall** > **NAT**.

Step 2 Choose **Add** > **Network Object NAT Rule**.

Step 3 Name the new network object, define the FTP server address, enable static NAT and enter the translated address.



Step 4 Click **Advanced** and configure the real and mapped interfaces and DNS modification.



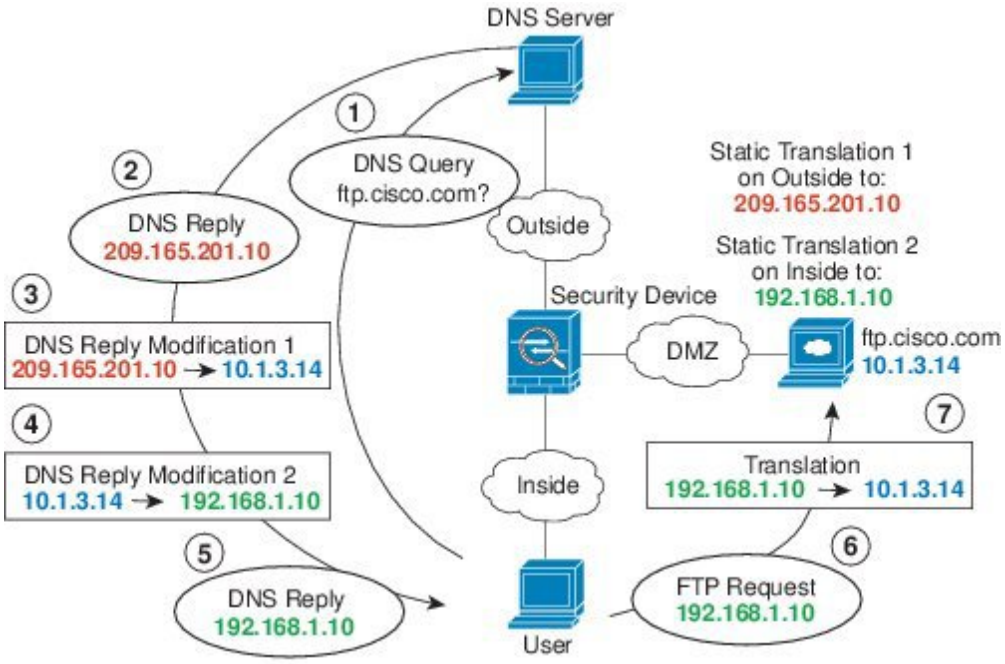
Step 5 Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

DNS Reply Modification, DNS Server, Host, and Server on Separate Networks

The following figure shows a user on the inside network requesting the IP address for ftp.cisco.com, which is on the DMZ network, from an outside DNS server. The DNS server replies with the mapped address (209.165.201.10) according to the static rule between outside and DMZ even though the user is not on the DMZ network. The ASA translates the address inside the DNS reply to 10.1.3.14.

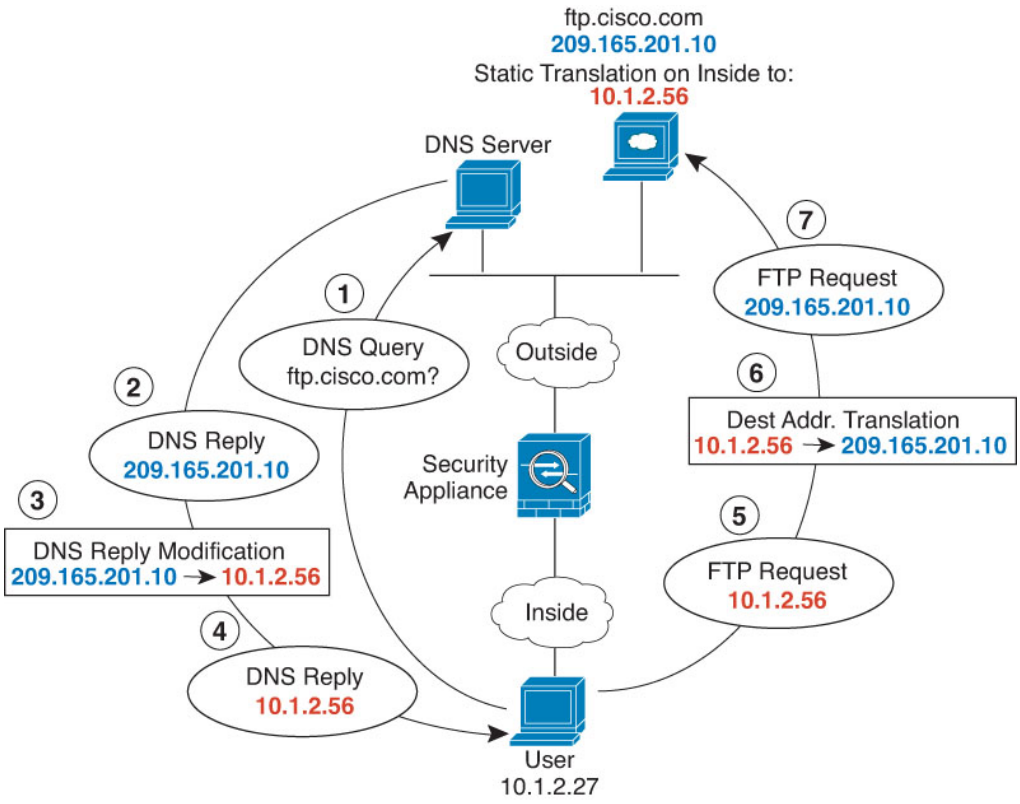
If the user needs to access ftp.cisco.com using the real address, then no further configuration is required. If there is also a static rule between the inside and DMZ, then you also need to enable DNS reply modification on this rule. The DNS reply will then be modified two times.In this case, the ASA again translates the address inside the DNS reply to 192.168.1.10 according to the static rule between inside and DMZ.

Figure 44: DNS Reply Modification, DNS Server, Host, and Server on Separate Networks



DNS Reply Modification, DNS Server on Host Network

The following figure shows an FTP server and DNS server on the outside. The system has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.20.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.

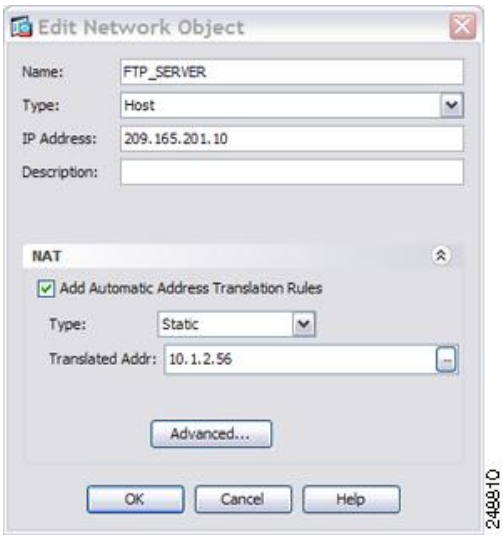


Procedure

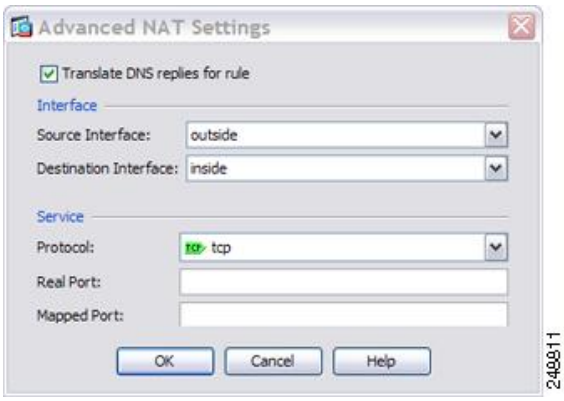
Step 1 Choose **Configuration** > **Firewall** > **NAT**.

Step 2 Choose **Add** > **Network Object NAT Rule**.

Step 3 Name the new network object, define the FTP server address, enable static NAT and enter the translated address.



Step 4 Click **Advanced** and configure the real and mapped interfaces and DNS modification.

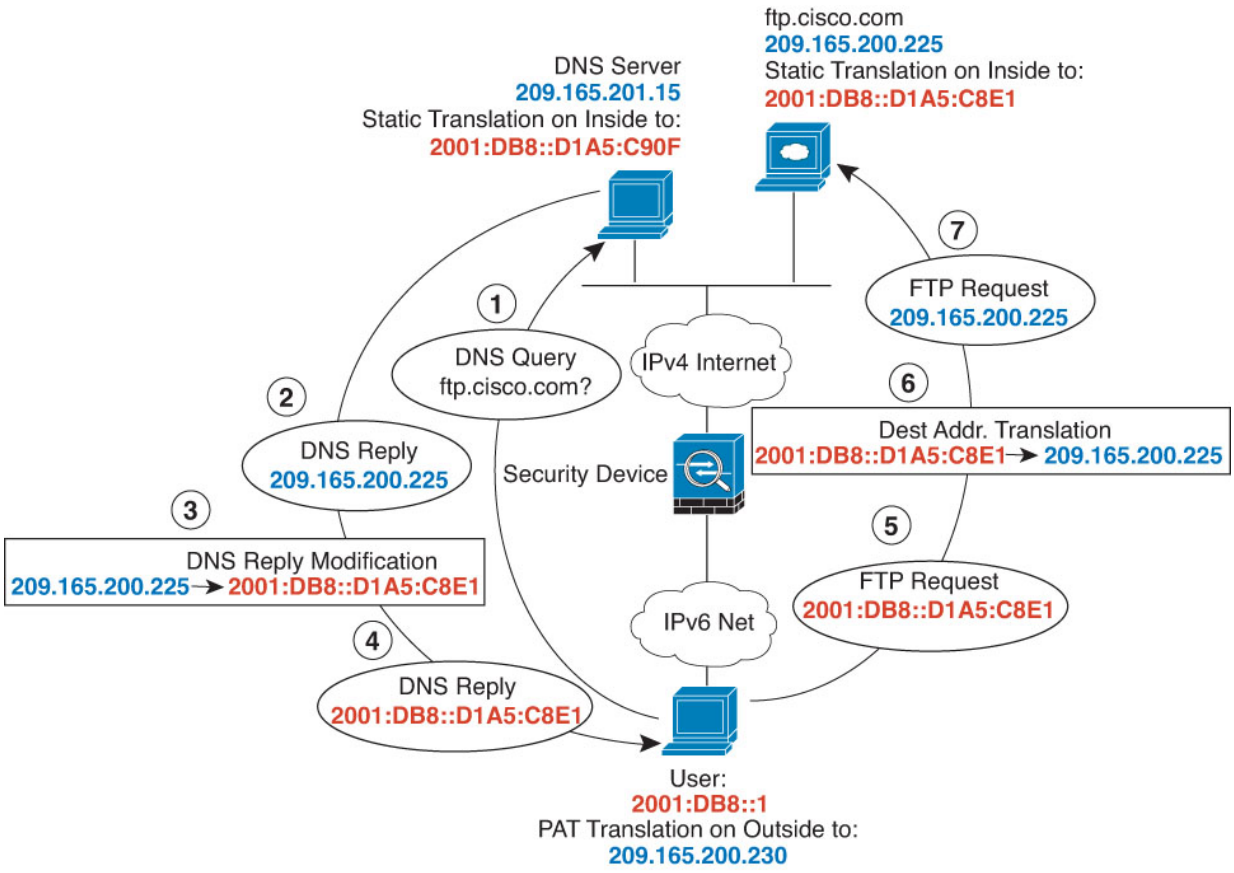


Step 5 Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

DNS64 Reply Modification

The following figure shows an FTP server and DNS server on the outside IPv4 network. The system has a static translation for the outside server. In this case, when an inside IPv6 user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.200.225.

Because you want inside users to use the mapped address for ftp.cisco.com (2001:DB8::D1A5:C8E1, where D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225) you need to configure DNS reply modification for the static translation. This example also includes a static NAT translation for the DNS server, and a PAT rule for the inside IPv6 hosts.



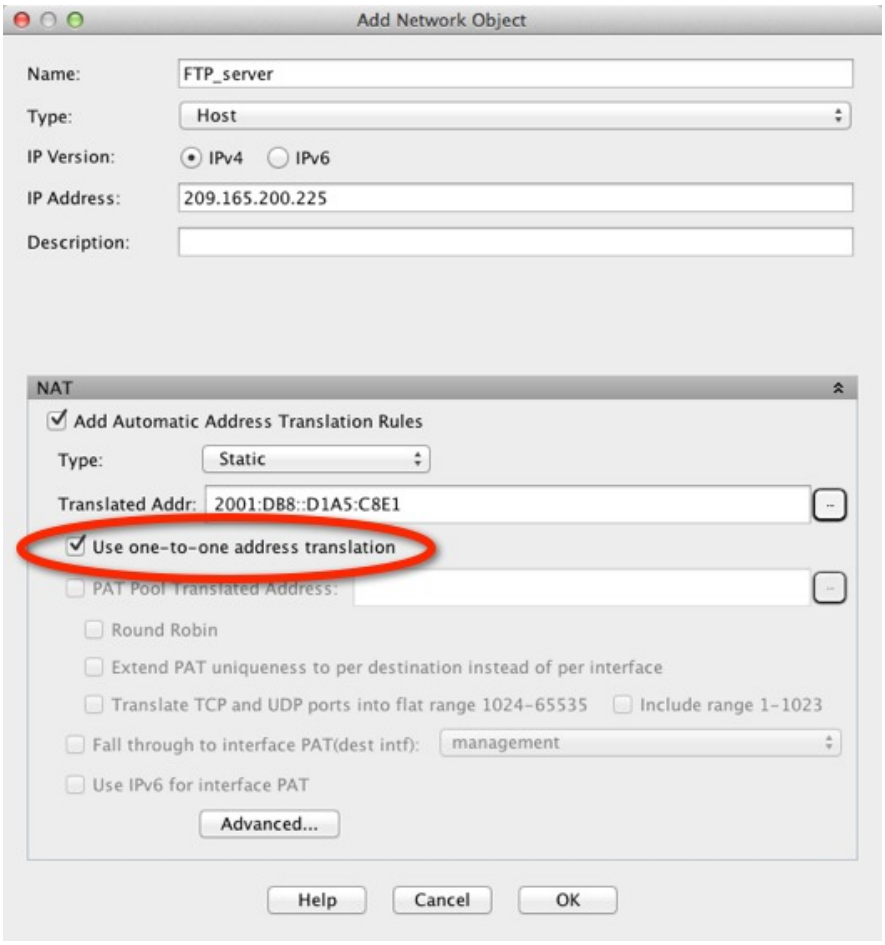
Procedure

Step 1 Choose **Configuration** > **Firewall** > **NAT**.

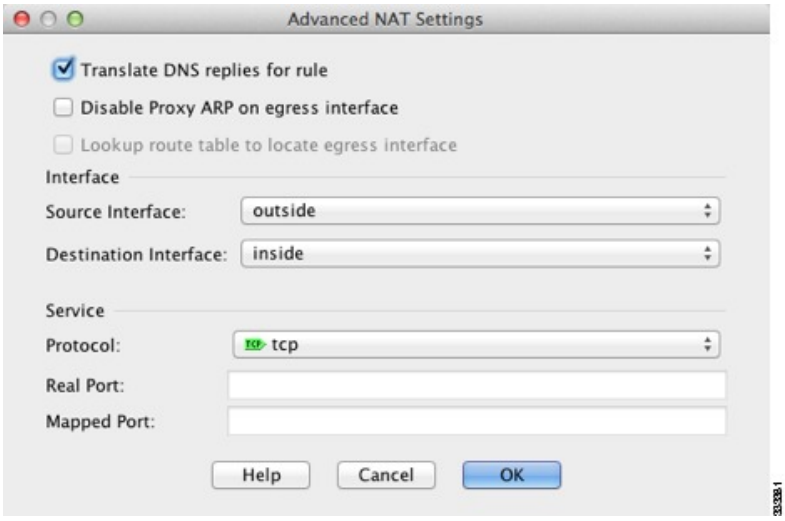
Step 2 Configure static network object NAT with DNS modification for the FTP server.

a) Choose **Add > Network Object NAT Rule**.

b) Name the new network object, define the FTP server address, enable static NAT, and enter the translated address. Because this is a one-to-one translation for NAT46, select **Use one-to-one address translation**.



c) Click **Advanced** to configure the real and mapped interfaces and DNS modification.

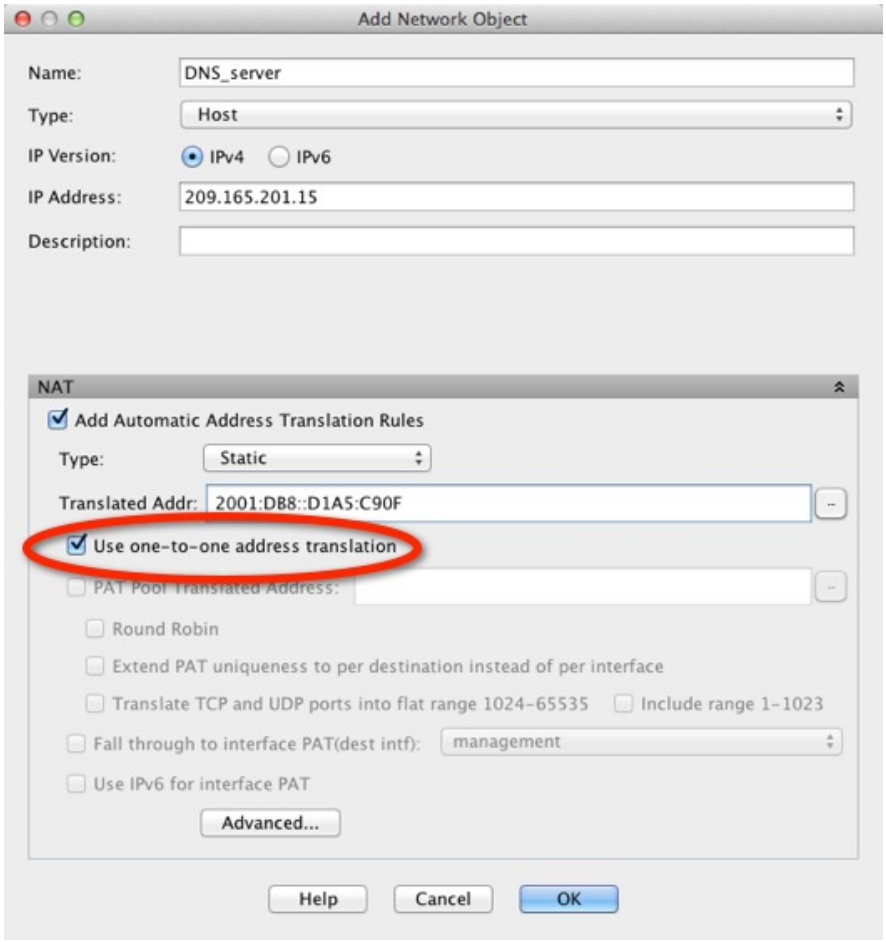


d) Click **OK** to return to the Network Object dialog box, and click **OK** again to save the rule.

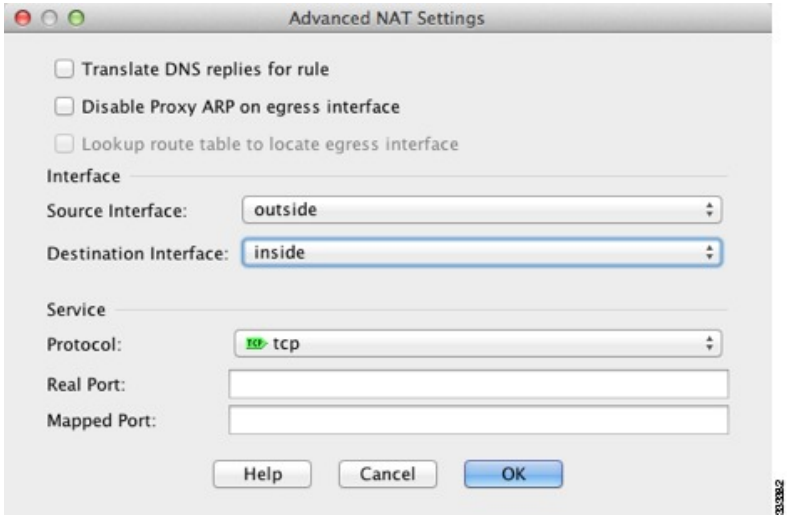
Step 3 Configure static network object NAT for the DNS server.

a) Choose **Add > Network Object NAT Rule**.

b) Name the new network object, define the DNS server address, enable static NAT, and enter the translated address. Because this is a one-to-one translation for NAT46, select **Use one-to-one address translation**.



c) Click **Advanced** to configure the real and mapped interfaces.



d) Click **OK** to return to the Network Object dialog box, and click **OK** again to save the rule.

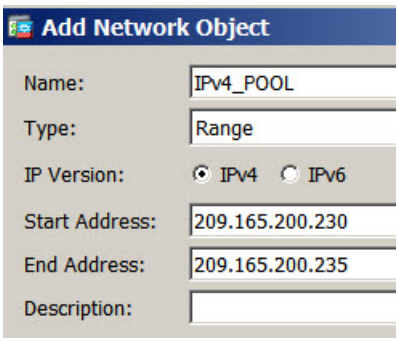
Step 4 Configure PAT for the inside IPv6 network.

a) Choose **Add > Network Object NAT Rule**.

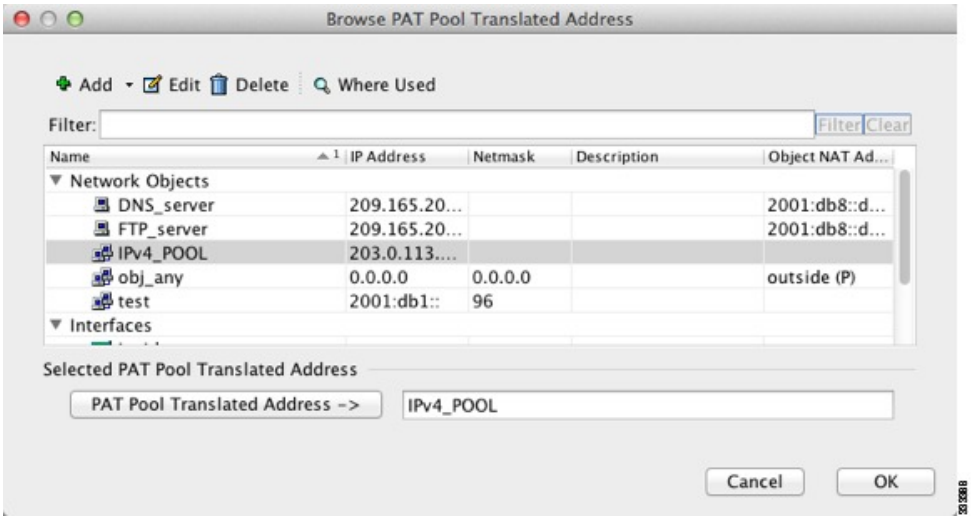
b) Name the new network object, define the IPv6 network address, and select **Dynamic** NAT.

c) Select **PAT Pool Translated Address**, and click the **...** (browse) button to create the PAT pool object.

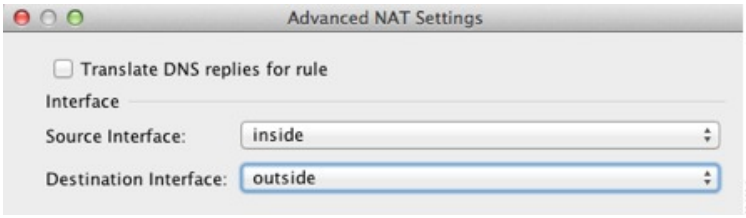
d) In the Browse PAT Pool Translated Address dialog box, select **Add > Network Object**. Name the new object, enter the address range for the PAT pool, and click **OK**.



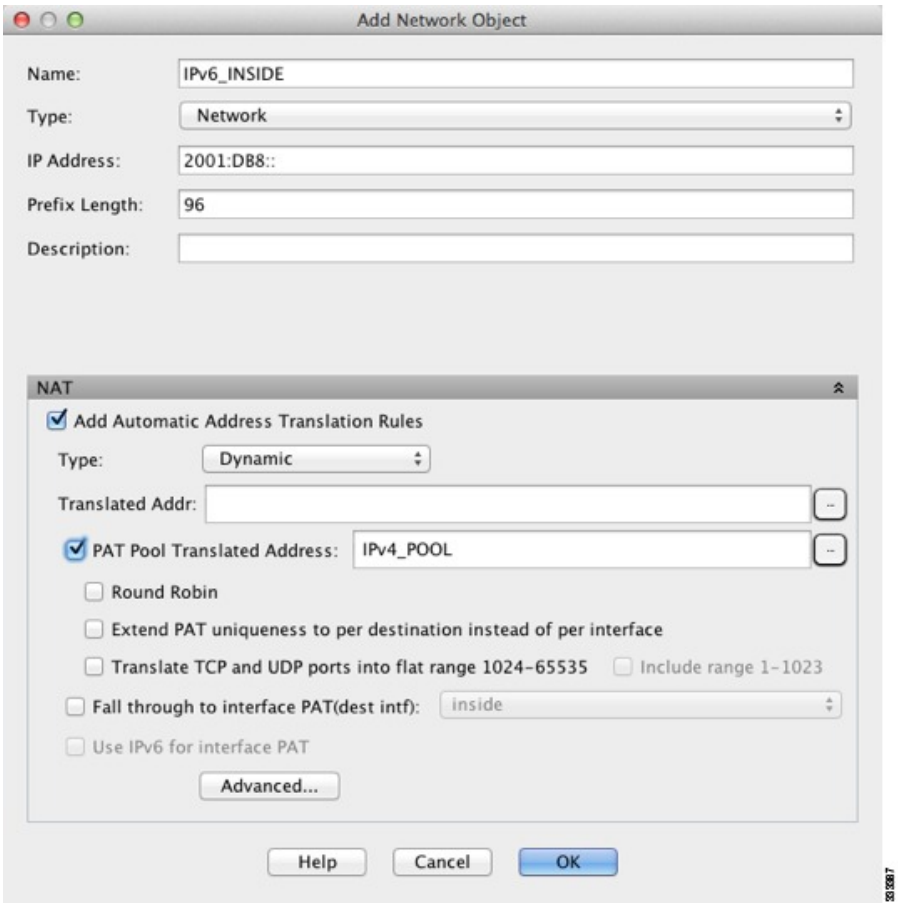
e) In the Browse PAT Pool Translated Address dialog box, double-click the PAT pool object you created to select it and click **OK**.



f) Click **Advanced** to configure the real and mapped interfaces.



g) Click **OK** to return to the Network Object dialog box.

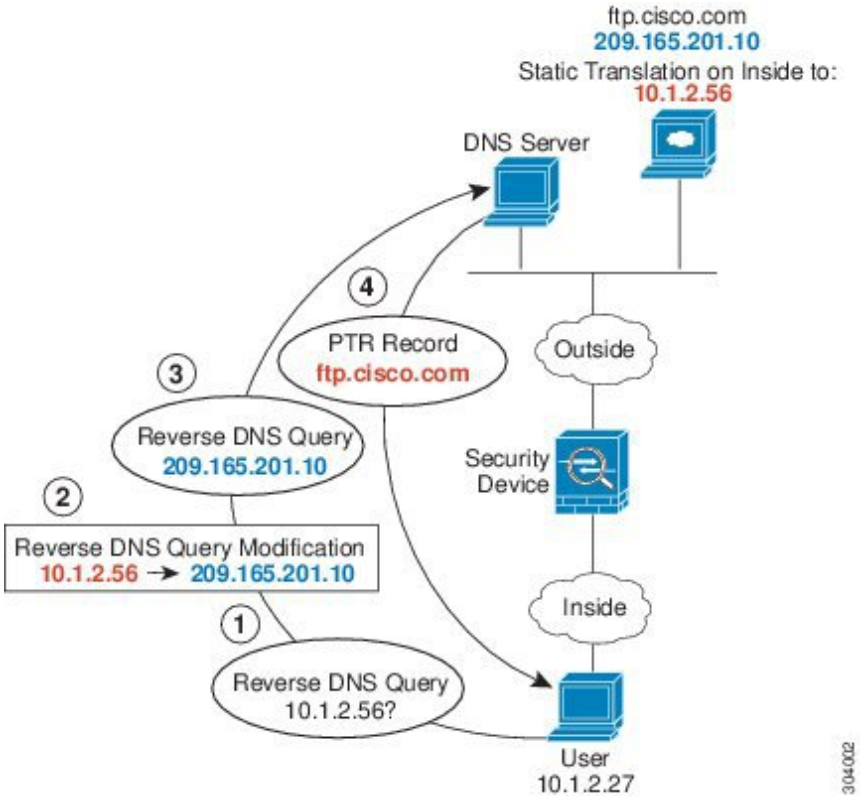


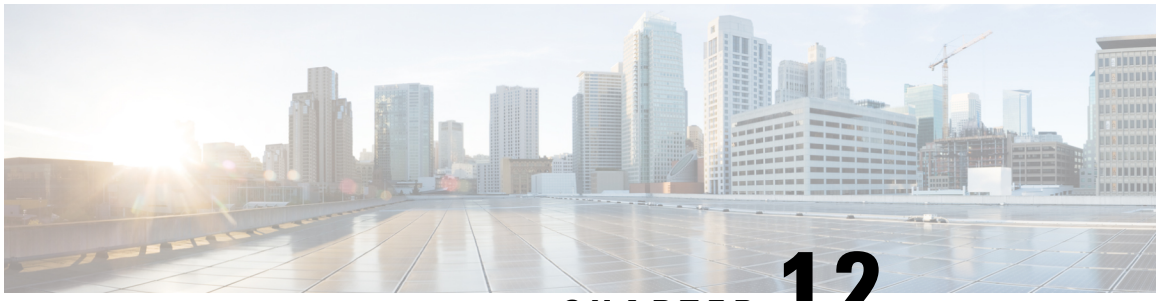
Step 5 Click **OK**, and then click **Apply**.

PTR Modification, DNS Server on Host Network

The following figure shows an FTP server and DNS server on the outside. The ASA has a static translation for the outside server. In this case, when an inside user performs a reverse DNS lookup for 10.1.2.56, the ASA modifies the reverse DNS query with the real address, and the DNS server responds with the server name, ftp.cisco.com.

Figure 45: PTR Modification, DNS Server on Host Network





CHAPTER 12

Mapping Address and Port (MAP)

Mapping Address and Port (MAP) is a carrier-grade feature for translating IPv4 addresses to IPv6, so the traffic can be sent over the service provider's IPv6 network before being translated back to IPv4 at the service provider edge.

- About Mapping Address and Port (MAP), on page 275
- Guidelines for Mapping Address and Port (MAP), on page 276
- Configure MAP-T Domains, on page 278
- Monitoring MAP, on page 279
- History for MAP, on page 280

About Mapping Address and Port (MAP)

Mapping Address and Port (MAP) is primarily a feature for use in service provider (SP) networks. The service provider can operate an IPv6-only network, the MAP domain, while supporting IPv4-only subscribers and their need to communicate with IPv4-only sites on the public Internet. MAP is defined in RFC7597, RFC7598, and RFC7599.

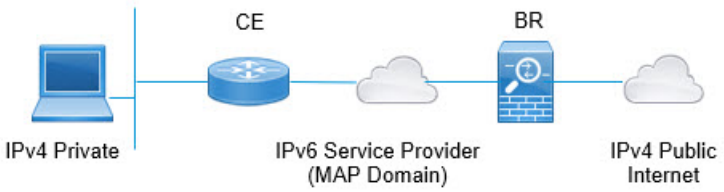
For the service provider, within the MAP domain, the benefit of MAP over NAT46 is that the substitution of an IPv6 address for the subscriber's IPv4 address (and back again to IPv4 at the SP network edge) is stateless. This provides greater efficiency within the SP network compared to NAT46.

There are two MAP techniques, MAP-Translation (MAP-T) and MAP-Encapsulation (MAP-E). The ASA supports MAP-T; MAP-E is not supported.

About Mapping Address and Port Translation (MAP-T)

With MAP-T, the subscriber's IPv4 address is first translated to the server provider's (SP) public IPv4 address, which could be either a one-to-one address mapping, or a mapping to a prefix or a shared address. Next, that IPv4 address is translated to an IPv6 address within the MAP domain, and the packet is transmitted over the SP IPv6 network. At the network edge, the SP's border relay is responsible for translating the IPv6 address back to the SP's IPv4 address before routing the packet to the public IPv4 network. The exact reverse is performed for traffic coming from the public IPv4 network to the subscriber.

Figure 46: MAP-T Network



By using MAP-T, you can transition the SP network to an IPv6-only architecture while allowed subscribers to continue using IPv4 and communicate with IPv4-only Internet or other sites outside the SP network.

MAP-T behaves like a NAT64 translation but instead of using an IPv6 address with an embedded IPv4 address, it uses an encoding scheme that also embeds the port number. Thus, MAP-T provides a way to restrict the port range used by devices.

A MAP-T system includes the following:

- **Customer Edge (CE) device**—The CE is a home gateway (wireless router, cable modem with router, and so forth). The CE provides IPv4/IPv6 translation as well as native IPv6 forwarding. It has one WAN-side provider-facing IPv6-addressed interface and one or more LAN-side interfaces addressed using private IPv4 addressing. You would configure one or more MAP domains for the CE to use to translate IPv4 packets to IPv6 and vice-verse.
- **Border Relay (BR) device**—You would install the ASA as a border relay. The BR is a provider-side component at the edge of the MAP domain that supports the IPv4/IPv6 translation. The BR has at least one IPv6-enabled interface and one IPv4 interface connected to the IPv4 network. You would configure one or more MAP domains for the BR to use to translate IPv4 packets to IPv6 and vice-verse. You must configure the CEs and BR with the same MAP domain rules.
- **MAP Domain**—A MAP domain is a mechanism to group a set of MAP-T CE devices with a set of MAP-T BR devices. A domain is a set of parameters that are shared between the BR and CE devices assigned to the domain. You configure the same domain with the same parameters on each of the BR and CE devices.

Guidelines for Mapping Address and Port (MAP)

Firewall Mode Guidelines

You can configure MAP in routed mode only. Transparent mode is not supported.

Additional Guidelines

- The ASA does not get involved in packet forwarding in mesh mode. Thus, you cannot configure forward mapping rules (FMR) in a MAP domain.
- MAP does not support tunneled VPN, multicast, or anycast traffic.
- You cannot use both NAT and MAP on a given connection. Ensure that your NAT and MAP rules do not overlap. You will get unexpected results if you have overlapping rules.
- The follow inspections do not support MAP translation. Packets subject to these inspections are not translated.

- CTIQBE
- DCERPC
- Diameter
- Name resolution over WINS
- GTP
- H.323 and H.225 and H.245 and RAS
- ILS (LDAP)
- Instant Messaging
- IP Options (RFC 791, 2113)
- IPSec Pass Through
- LISP
- M3UA
- MGCP
- MMP
- NetBIOS
- PPTP
- RADIUS accounting
- RSH
- RTSP
- SIP
- SKINNY
- SMTP and ESMTP
- SNMP
- SQL*Net
- STUN
- Sun RPC
- TFTP
- WAAS
- XDMCP
- Active FTP

Configure MAP-T Domains

To configure MAP-T, you create one or more domains. When you configure MAP-T on customer edge (CE) and border relay (BR) devices, ensure that you use the same parameters for each device that will participate in each domain.

You can configure up to 25 MAP-T domains. In multiple-context mode, you can configure up to 25 domains per context.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **CGNAT Map**.

Step 2 Do any of the following:

- Click **Add** to create a new MAP domain.
- Select a MAP domain and click **Edit** to modify the domain.

If you no longer need the domain, select it and click **Delete**.

Step 3 In **MAP Domain Name**, enter a name for the domain. The name is an alphanumeric string up to 48 characters. The name can also include the following special characters: period (.), slash (/), and colon (:).

Step 4 Click the **Default Mapping Rule** tab and configure the **Rule IPV6 Prefix** and **Rule IPv6 Prefix Length** for the rule.

Specify an IPv6 prefix to be used to embed IPv4 destination addresses per RFC 6052. The prefix length should normally be 64, but allowed values are 32, 40, 48, 56, 64 or 96. Any trailing bits after the embedded IPv4 address are set to 0. For example, 2001:DB8:CAFE:CAFE::/64.

The border relay (BR) device uses this rule to translate all IPv4 addresses outside the MAP domain to an IPv6 address that works within the MAP domain.

Step 5 Click the **Basic Mapping Rule** tab, configure the basic mapping rule IP address prefixes and port parameters.

The customer edge (CE) device uses the basic mapping rule to determine its dedicated IPv4 addressing or shared address and port set assignment. The CE device first translates the system's IPv4 address to an IPv4 address and port within the pool's prefix and port range (using NAT44), then MAP translates the new IPv4 address to an IPv6 address within the pool defined by the rule's IPv6 prefix. The packet is then ready to be transmitted over the service provider's IPv6-only network to a border relay (BR) device.

Configure the following options:

- **Rule IPV4 Prefix**, **Rule IPv4 Subnet Mask**—The IPv4 prefix defines the IPv4 address pool for the customer edge (CE) device. The CE device first translates its IPv4 address to an address (and port number) in the pool defined by the IPv4 prefix. MAP then translates this new address to an IPv6 address using the prefix in the default mapping rule.

  Specify a network address and subnet mask, for example, 192.168.3.0 255.255.255.0. You cannot use the same IPv4 prefix in different MAP domains.

- **Rule IPV6 Prefix**, **Rule IPv6 Prefix Length**—The IPv6 prefix defines the address pool for the CE device's IPv6 address. MAP translates IPv6 packets back to IPv4 only if the packets have a destination address with this prefix and a source address with the IPv6 prefix defined in the default mapping rule,

and is within the right port range. Any IPv6 packets sent to the CE device from other addresses are simply processed as IPv6 traffic without MAP translation. Packets from the MAP source/destination pools, but with out-of-range ports, are simply dropped.

Specify an IPv6 prefix and prefix length, which is normally 64, but cannot be less than 8. You cannot use the same IPv6 prefix in different MAP domains. For example, 2001:DB8:FFFF:F000::/64.

- **Share Ratio**—Specify the number of ports that should be in the pool. The number must be a power of 2, from 1-65536, such as 1, 2, 4, 8, and so forth.
- **Start Port**—The first port in the port pool for the translated address. The port you specify must be a power of 2, from 1-32768 such as 1, 2, 4, 8, and so forth. If you want to exclude the well-known ports, start at 1024 or higher.

Step 6 Click **OK**.

Monitoring MAP

The following topics explain how you can monitor the MAP configuration and activity.

Verifying the MAP Domain Configuration

You can view the map domains and their status to verify the configuration is correct.

Select **Monitoring** > **Properties**, then select **MAP Domains** from the table of contents. The information includes the MAP configuration, and shows the output of the **show map-domain** command. If the configuration for a domain is not yet complete, this is indicated. An incompletely configured domain is not active. You can enter a map name and click **Filter** to see information for a single domain.

```
MAP Domain 1
  Default Mapping Rule
    IPv6 prefix 2001:db8:cafe:cafe::/64
  Basic Mapping Rule
    IPv6 prefix 2001:cafe:cafe:1::/64
    IPv4 prefix 192.168.3.0 255.255.255.0
    share ratio 16
    start port 1024
    PSID length 4
    PSID offset 6
    Rule EA-bit length 12

MAP Domain 2
  Default Mapping Rule
    IPv6 prefix 2001:db8:1234:1234::/64

Warning: map-domain 2 configuration is incomplete and not in effect.
```

Monitoring MAP Syslog Messages

If you enable syslog, you can monitor MAP behavior with the following syslog messages:

- 305018: MAP translation from *interface name:source IP address/source port-destination IP address/destination port* to *interface name:translated source IP address/translated source port-translated destination IP address/translated destination port*

  A new MAP translation was made. The message shows the original and translated source and destination.

- 305019: MAP node address *IP address/port* has inconsistent Port Set ID encoding

  A packet has an address that matches MAP basic mapping rules, which means it is meant to be translated, but the Port Set ID encoded within the address is inconsistent per RFC7599. This likely means there is a software fault on the MAP node where this packet originates.

- 305020: MAP node with address *IP address* is not allowed to use port *port*

  A packet has an address that matches MAP basic mapping rules, which means it is meant to be translated, but the associated port does not fall within the range allocated to that address. This likely means there is misconfiguration on the MAP node where this packet originates.

History for MAP

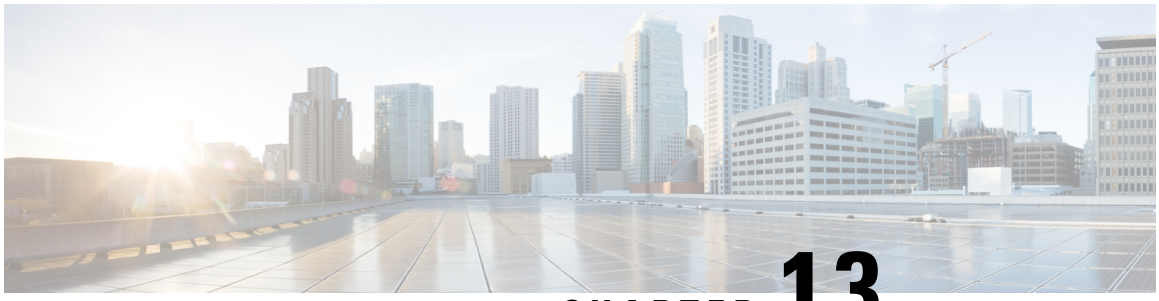
Feature Name	Platform Releases	Description
Mapping Address and Port-Translation (MAP-T)	9.13(1)	Mapping Address and Port (MAP) is primarily a feature for use in service provider (SP) networks. The service provider can operate an IPv6-only network, the MAP domain, while supporting IPv4-only subscribers and their need to communicate with IPv4-only sites on the public Internet. MAP is defined in RFC7597, RFC7598, and RFC7599. We introduced or modified the following screens: **Configuration** > **Device Setup** > **CGNAT Map**, **Monitoring** > **Properties** > **MAP Domains**.



PART IV

Service Policies and Application Inspection

- Service Policy, on page 283
- Getting Started with Application Layer Protocol Inspection, on page 299
- Inspection of Basic Internet Protocols, on page 317
- Inspection for Voice and Video Protocols, on page 351
- Inspection for Mobile Networks, on page 373



CHAPTER 13

Service Policy

Service policies provide a consistent and flexible way to configure ASA features. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. A service policy consists of multiple actions or rules applied to an interface or applied globally.

- About Service Policies, on page 283
- Guidelines for Service Policies, on page 289
- Defaults for Service Policies, on page 290
- Configure Service Policies, on page 291
- History for Service Policies, on page 297

About Service Policies

The following topics describe how service policies work.

The Components of a Service Policy

The point of service policies is to apply advanced services to the traffic you are allowing. Any traffic permitted by access rules can have service policies applied, and thus receive special processing, such as being redirected to a service module or having application inspection applied.

You can have these types of service policy:

- One global policy that gets applied to all interfaces.
- One service policy applied per interface. The policy can be a mix of classes for traffic going through the device and management traffic directed at the ASA interface rather than going through it,

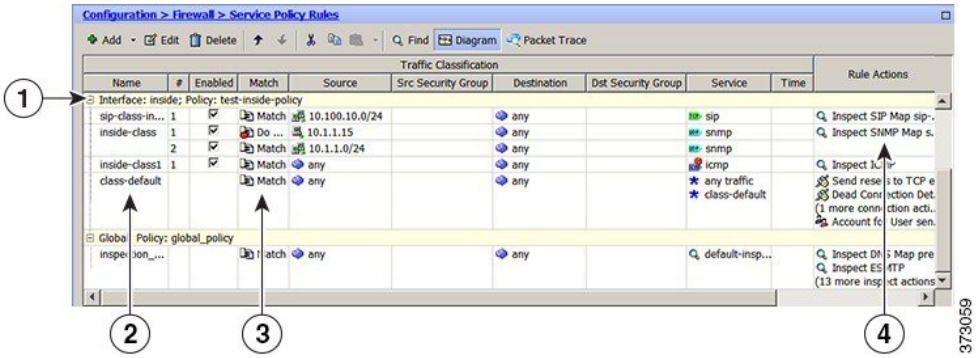
Each service policy is composed of the following elements:

1. Service policy map, which is the ordered set of rules, and is named on the **service-policy** command. In ASDM, the policy map is represented as a folder on the Service Policy Rules page.
2. Rules, each rule being a **class** command within the service policy map and the commands associated with the **class** command. In ASDM, each rule is shown on a separate row, and the name of the rule is the class name.

   The **class** command defines the traffic matching criteria for the rule.

The commands associated with class, such as **inspect**, **set connection timeout**, and so forth, define the services and constraints to apply to matching traffic. Note that inspect commands can point to inspection policy maps, which define actions to apply to inspected traffic. Keep in mind that inspection policy maps are not the same as service policy maps.

The following example compares how service policies appear in the CLI with how they appear in ASDM. Note that there is not a one-to-one mapping between the figure call-outs and lines in the CLI.



The following CLI is generated by the rules shown in the figure above.

```
: Access lists used in class maps.
: In ASDM, these map to call-out 3, from the Match to the Time fields.
access-list inside_mpc line 1 extended permit tcp 10.100.10.0 255.255.255.0 any eq sip
access-list inside_mpc_1 line 1 extended deny udp host 10.1.1.15 any eq snmp
access-list inside_mpc_1 line 2 extended permit udp 10.1.1.0 255.255.255.0 any eq snmp
access-list inside_mpc_2 line 1 extended permit icmp any any
: SNMP map for SNMP inspection. Denies all but v3.
: In ASDM, this maps to call-out 4, rule actions, for the class-inside policy.
snmp-map snmp-v3only
  deny version 1
  deny version 2
  deny version 2c
: Inspection policy map to define SIP behavior.
: The sip-high inspection policy map must be referred to by an inspect sip command
: in the service policy map.
: In ASDM, this maps to call-out 4, rule actions, for the sip-class-inside policy.
policy-map type inspect sip sip-high
  parameters
    rtp-conformance enforce-payloadtype
    no traffic-non-sip
    software-version action mask log
    uri-non-sip action mask log
    state-checking action drop-connection log
    max-forwards-validation action drop log
    strict-header-validation action drop log
: Class map to define traffic matching for the inside-class rule.
: In ASDM, this maps to call-out 3, from the Match to the Time fields.
class-map inside-class
  match access-list inside_mpc_1
: Class map to define traffic matching for the sip-class-inside rule.
: In ASDM, this maps to call-out 3, from the Match to the Time fields.
class-map sip-class-inside
  match access-list inside_mpc
: Class map to define traffic matching for the inside-class1 rule.
: In ASDM, this maps to call-out 3, from the Match to the Time fields.
class-map inside-class1
  match access-list inside_mpc_2
```

```
: Policy map that actually defines the service policy rule set named test-inside-policy.
: In ASDM, this corresponds to the folder at call-out 1.
policy-map test-inside-policy
: First rule in test-inside-policy, named sip-class-inside. Inspects SIP traffic.
: The sip-class-inside rule applies the sip-high inspection policy map to SIP inspection.
: In ASDM, each rule corresponds to call-out 2.
  class sip-class-inside
    inspect sip sip-high
: Second rule, inside-class. Applies SNMP inspection using an SNMP map.
  class inside-class
    inspect snmp snmp-v3only
: Third rule, inside-class1. Applies ICMP inspection.
  class inside-class1
    inspect icmp
: Fourth rule, class-default. Applies connection settings and enables user statistics.
  class class-default
    set connection timeout embryonic 0:00:30 half-closed 0:10:00 idle 1:00:00
reset dcd 0:15:00 5
    user-statistics accounting
: The service-policy command applies the policy map rule set to the inside interface.
: This command activates the policies.
service-policy test-inside-policy interface inside
```

Features Configured with Service Policies

The following table lists the features you configure using service policies.

Table 10: Features Configured with Service Policies

Feature	For Through Traffic?	For Management Traffic?	See:
Application inspection (multiple types)	**All** except RADIUS accounting	**RADIUS accounting** only	• Getting Started with Application Layer Protocol Inspection, on page 299. • Inspection of Basic Internet Protocols, on page 317. • Inspection for Voice and Video Protocols, on page 351. • Inspection for Mobile Networks, on page 373.
ASA FirePOWER (ASA SFR)	**Yes**	No	ASA FirePOWER Module, on page 97.
NetFlow Secure Event Logging filtering	**Yes**	**Yes**	See the NetFlow implementation guide.
QoS input and output policing	**Yes**	No	Quality of Service, on page 437.
QoS standard priority queue	**Yes**	No	Quality of Service, on page 437.
TCP and UDP connection limits and timeouts, and TCP sequence number randomization	**Yes**	**Yes**	Connection Settings, on page 413.
TCP normalization	**Yes**	No	Connection Settings, on page 413.

Feature	For Through Traffic?	For Management Traffic?	See:
TCP state bypass	**Yes**	No	Connection Settings, on page 413.
User statistics for Identity Firewall	**Yes**	**Yes**	See the **user-statistics** command in the command reference.

Feature Directionality

Actions are applied to traffic bidirectionally or unidirectionally depending on the feature. For features that are applied bidirectionally, all traffic that enters or exits the interface to which you apply the policy map is affected if the traffic matches the class map for both directions.



Note When you use a global policy, all features are unidirectional; features that are normally bidirectional when applied to a single interface only apply to the ingress of each interface when applied globally. Because the policy is applied to all interfaces, the policy will be applied in both directions so bidirectionality in this case is redundant.

For features that are applied unidirectionally, for example QoS priority queue, only traffic that enters (or exits, depending on the feature) the interface to which you apply the policy map is affected. See the following table for the directionality of each feature.

Table 11: Feature Directionality

Feature	Single Interface Direction	Global Direction
Application inspection (multiple types)	Bidirectional	Ingress
ASA FirePOWER (ASA SFR)	Bidirectional	Ingress
NetFlow Secure Event Logging filtering	N/A	Ingress
QoS input policing	Ingress	Ingress
QoS output policing	Egress	Egress
QoS standard priority queue	Egress	Egress
TCP and UDP connection limits and timeouts, and TCP sequence number randomization	Bidirectional	Ingress
TCP normalization	Bidirectional	Ingress
TCP state bypass	Bidirectional	Ingress
User statistics for Identity Firewall	Bidirectional	Ingress

Feature Matching Within a Service Policy

A packet matches rules in a policy for a given interface according to the following rules:

1. A packet can match only one rule for an interface for each feature type.
2. When the packet matches a rule for a feature type, the ASA does not attempt to match it to any subsequent rules for that feature type.
3. If the packet matches a subsequent rule for a different feature type, however, then the ASA also applies the actions for the subsequent rule, if supported. See Incompatibility of Certain Feature Actions, on page 288 for more information about unsupported combinations.

Note Application inspection includes multiple inspection types, and most are mutually exclusive. For inspections that can be combined, each inspection is considered to be a separate feature.

Examples of Packet Matching

For example:

- If a packet matches a rule for connection limits, and also matches a rule for an application inspection, then both actions are applied.
- If a packet matches a rule for HTTP inspection, but also matches another rule that includes HTTP inspection, then the second rule actions are not applied.
- If a packet matches a rule for HTTP inspection, but also matches another rule that includes FTP inspection, then the second rule actions are not applied because HTTP and FTP inspections cannot be combined.
- If a packet matches a rule for HTTP inspection, but also matches another rule that includes IPv6 inspection, then both actions are applied because the IPv6 inspection can be combined with any other type of inspection.

Order in Which Multiple Feature Actions are Applied

The order in which different types of actions in a service policy are performed is independent of the order in which the actions appear in the table.

Actions are performed in the following order:

1. QoS input policing
2. TCP normalization, TCP and UDP connection limits and timeouts, TCP sequence number randomization, and TCP state bypass.

Note When a the ASA performs a proxy service (such as AAA) or it modifies the TCP payload (such as FTP inspection), the TCP normalizer acts in dual mode, where it is applied before and after the proxy or payload modifying service.

3. Application inspections that can be combined with other inspections:
   a. IPv6
   b. IP options
   c. WAAS
4. Application inspections that cannot be combined with other inspections. See Incompatibility of Certain Feature Actions, on page 288 for more information.
5. ASA FirePOWER (ASA SFR)
6. QoS output policing
7. QoS standard priority queue

Note NetFlow Secure Event Logging filtering and User statistics for Identity Firewall are order-independent.

Incompatibility of Certain Feature Actions

Some features are not compatible with each other for the same traffic. The following list might not include all incompatibilities; for information about compatibility of each feature, see the chapter or section for the feature:

- You cannot configure QoS priority queuing and QoS policing for the same set of traffic.
- Most inspections should not be combined with another inspection, so the ASA only applies one inspection if you configure multiple inspections for the same traffic. Exceptions are listed in Order in Which Multiple Feature Actions are Applied, on page 287.
- You cannot configure traffic to be sent to multiple modules.
- HTTP inspection is not compatible with ASA FirePOWER.

Note The Default Inspection Traffic traffic class, which is used in the default global policy, is a special CLI shortcut to match the default ports for all inspections. When used in a policy map, this class map ensures that the correct inspection is applied to each packet, based on the destination port of the traffic. For example, when UDP traffic for port 69 reaches the ASA, then the ASA applies the TFTP inspection; when TCP traffic for port 21 arrives, then the ASA applies the FTP inspection. So in this case only, you can configure multiple inspections for the same class map. Normally, the ASA does not use the port number to determine which inspection to apply, thus giving you the flexibility to apply inspections to non-standard ports, for example.

Feature Matching for Multiple Service Policies

For TCP and UDP traffic (and ICMP when you enable stateful ICMP inspection), service policies operate on traffic flows, and not just individual packets. If traffic is part of an existing connection that matches a feature

in a policy on one interface, that traffic flow cannot also match the same feature in a policy on another interface; only the first policy is used.

For example, if HTTP traffic matches a policy on the inside interface to inspect HTTP traffic, and you have a separate policy on the outside interface for HTTP inspection, then that traffic is not also inspected on the egress of the outside interface. Similarly, the return traffic for that connection will not be inspected by the ingress policy of the outside interface, nor by the egress policy of the inside interface.

For traffic that is not treated as a flow, for example ICMP when you do not enable stateful ICMP inspection, returning traffic can match a different policy map on the returning interface.

Guidelines for Service Policies

Inspection Guidelines

There is a separate topic that provides detailed guidelines for application inspection service policies. See Guidelines for Application Inspection, on page 301.

IPv6 Guidelines

Supports IPv6 for the following features:

- Application inspection for several, but not all, protocols. For details, see Guidelines for Application Inspection, on page 301.
- ASA FirePOWER
- NetFlow Secure Event Logging filtering
- SCTP state bypass
- TCP and UDP connection limits and timeouts, TCP sequence number randomization
- TCP normalization
- TCP state bypass
- User statistics for Identity Firewall

Class Map (Traffic Class) Guidelines

The maximum number of class maps (traffic classes) of all types is 255 in single mode or per context in multiple mode. Class maps include the following types:

- Layer 3/4 class maps (for through traffic and management traffic).
- Inspection class maps
- Regular expression class maps
- **match** commands used directly underneath an inspection policy map

This limit also includes default class maps of all types, limiting user-configured class maps to approximately 235.

Service Policy Guidelines

- Interface service policies on ingress interfaces take precedence over the global service policy for a given feature. For example, if you have a global policy with FTP inspection, and an interface policy with TCP normalization, then both FTP inspection and TCP normalization are applied to the interface. However, if you have a global policy with FTP inspection, and an ingress interface policy with FTP inspection, then only the ingress interface policy FTP inspection is applied to that interface. If no ingress or global policy implements a feature, then an interface service policy on the egress interface that specifies the feature is applied.
- You can only apply one global policy. For example, you cannot create a global policy that includes feature set 1, and a separate global policy that includes feature set 2. All features must be included in a single policy.
- When you make service policy changes to the configuration, all *new* connections use the new service policy. Existing connections continue to use the policy that was configured at the time of the connection establishment. Output for the **show** command will not include data about the old connections.

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

  To ensure that all connections use the new policy, you need to disconnect the current connections so they can reconnect using the new policy. Use the **clear conn** or **clear local-host** commands.

Defaults for Service Policies

The following topics describe the default settings for service policies and the Modular Policy Framework.

Default Service Policy Configuration

By default, the configuration includes a policy that matches all default application inspection traffic and applies certain inspections to the traffic on all interfaces (a global policy). Not all inspections are enabled by default. You can only apply one global policy, so if you want to alter the global policy, you need to either edit the default policy or disable it and apply a new one. (An interface policy overrides the global policy for a particular feature.)

The default policy includes the following application inspections:

- DNS
- FTP
- H323 (H225)
- H323 (RAS)
- RSH
- RTSP
- ESMTP
- SQLnet

- Skinny (SCCP)
- SunRPC
- SIP
- NetBios
- TFTP
- IP Options

Default Class Maps (Traffic Classes)

The configuration includes a default Layer 3/4 class map (traffic class) that the ASA uses in the default global policy called Default Inspection Traffic; it matches the default inspection traffic. This class, which is used in the default global policy, is a special shortcut to match the default ports for all inspections.

When used in a policy, this class ensures that the correct inspection is applied to each packet, based on the destination port of the traffic. For example, when UDP traffic for port 69 reaches the ASA, then the ASA applies the TFTP inspection; when TCP traffic for port 21 arrives, then the ASA applies the FTP inspection. So in this case only, you can configure multiple inspections for the same class map. Normally, the ASA does not use the port number to determine which inspection to apply, thus giving you the flexibility to apply inspections to non-standard ports, for example.

Another class map that exists in the default configuration is called class-default, and it matches all traffic. This class map appears at the end of all Layer 3/4 policy maps and essentially tells the ASA to not perform any actions on all other traffic. If you use the class-default class (the Any traffic class in ASDM), use it for broader purposes like enabling TCP sequence randomization, TTL decrement, or any other optional attribute changes. Do not use it for inspection, or it can cause traffic disruptions. Some features are only available for class-default.

Configure Service Policies

Configuring a service policy consists of adding one or more service policy rules per interface or for the global policy. ASDM uses a wizard to take you through the process of creating a service policy. For each rule, you identify the following elements:

1. The interface to which you want to apply the rule, or the global policy.
2. The traffic to which you want to apply actions. You can identify Layer 3 and 4 traffic.
3. The actions to apply to the traffic class. You can apply multiple non-conflicting actions for each traffic class.

After you create a policy, you can add rules, move, edit, or delete rules or policies. The following topics explain how to configure service policies.

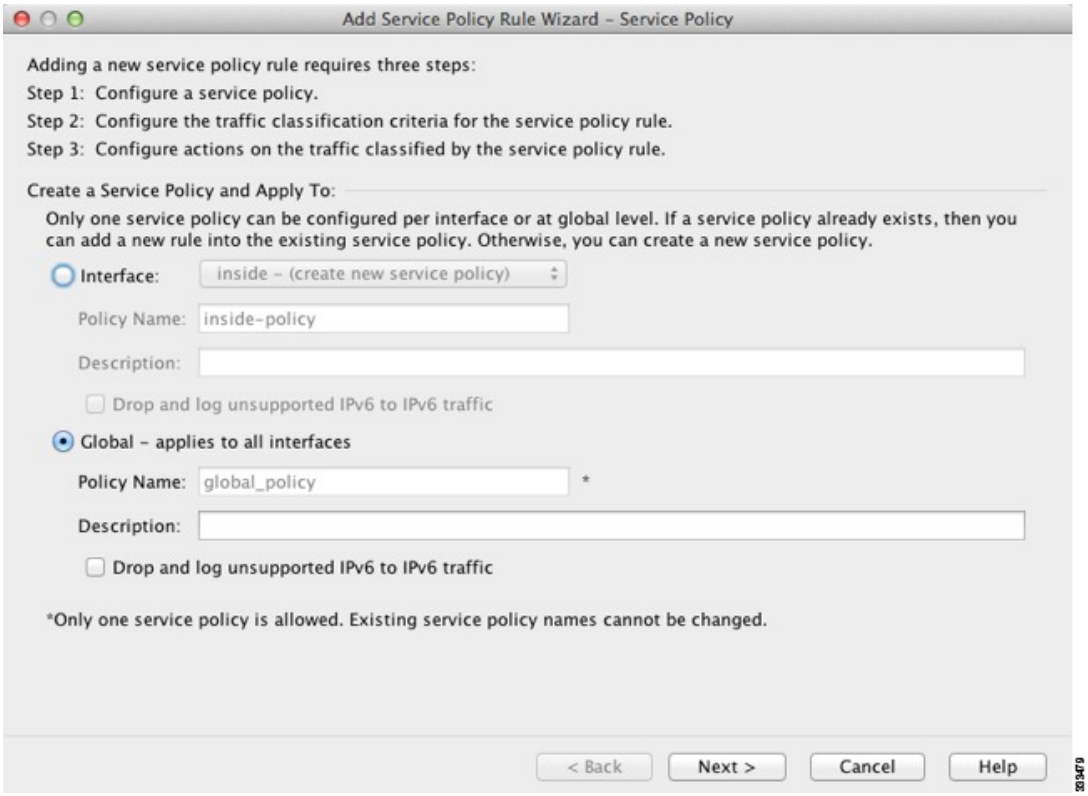
Add a Service Policy Rule for Through Traffic

To add a service policy rule for through traffic, use the Add Service Policy Rule wizard. You will be asked to choose the scope of the policy, for a specific interface or global:

- Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with FTP inspection, and an interface policy with TCP connection limits, then both FTP inspection and TCP connection limits are applied to the interface. However, if you have a global policy with FTP inspection, and an interface policy with FTP inspection, then only the interface policy FTP inspection is applied to that interface.
- Global service policies provide default services to all interfaces. Unless overridden by an interface-specific policy, the global services are applied. By default, a global policy exists that includes a service policy rule for default application inspection. You can add a rule to the global policy using the wizard.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy Rules**, and click **Add** or **Add** > **Add Service Policy Rule**.



Step 2 In the Create a Service Policy and Apply To area:

a) Choose whether the policy applies to a specific **Interface** or **Global** to all interfaces.
b) If you select Interface, choose the name of the interface. If the interface already has a policy, then you are adding a rule to the existing policy.
c) If the interface does not already have a service policy, enter the name of the new policy.
d) (Optional) Enter a description for the policy.

e) (Optional) Check the **Drop and log unsupported IPv6 to IPv6 traffic** option to generate a syslog (767001) for IPv6 traffic that is dropped by application inspections that do not support IPv6 traffic. By default, syslogs are not generated.

f) Click **Next**.

Step 3 On the Traffic Classification Criteria page, choose one of the following options to specify the traffic to which to apply the policy actions and click **Next**.

- **Create a new traffic class**. Enter a traffic class name and an optional description.

  Identify the traffic using one of several criteria:

  - **Default Inspection Traffic**—The class matches the default TCP and UDP ports used by all applications that the ASA can inspect. When you click **Next**, you are shown the services and ports defined by this class.

    This option, which is used in the default global policy, is a special shortcut that when used in a rule, ensures that the correct inspection is applied to each packet, based on the destination port of the traffic. For more information, see Default Class Maps (Traffic Classes), on page 291.

    See Default Inspections and NAT Limitations, on page 302 for a list of default ports. The ASA includes a default global policy that matches the default inspection traffic, and applies common inspections to the traffic on all interfaces. Not all applications whose ports are included in the Default Inspection Traffic class are enabled by default in the policy map.

    You can specify a Source and Destination IP Address class (which uses an ACL) along with the Default Inspection Traffic class to narrow the matched traffic. Because the Default Inspection Traffic class specifies the ports and protocols to match, any ports and protocols in the ACL are ignored.

  - **Source and Destination IP Address (uses ACL)**—The class matches traffic specified by an extended ACL. When you click **Next**, you are prompted for the attributes of the access control entry, and the wizard builds the ACL. Optionally, you can select an existing ACL.

    When defining the ACE, the Match option creates a rule where traffic matching the addresses have actions applied. The Do Not Match option exempts the traffic from having the specified actions applied. For example, you want to match all traffic in 10.1.1.0/24 and apply connection limits to it, except for 10.1.1.25. In this case, create two rules, one for 10.1.1.0/24 using the Match option and one for 10.1.1.25 using the Do Not Match option. Be sure to arrange the rules so that the Do Not Match rule is above the Match rule, or else 10.1.1.25 will match the Match rule first.

    Note
    When you create a new traffic class of this type, you can only specify one access control entry (ACE) initially. After you finish adding the rule, you can add additional ACEs by adding a new rule to the same interface or global policy, and then specifying **Add rule to existing traffic class** (see below).

  - **Tunnel Group**—The class matches traffic for a tunnel group (connection profile) to which you want to apply QoS. You can also specify one other traffic match option to refine the traffic match, excluding Any Traffic, Source and Destination IP Address (uses ACL), or Default Inspection Traffic.

    When you click **Next**, you are prompted to select the tunnel group (you can create a new one if necessary). To police each flow, check **Match flow destination IP address**. All traffic going to a unique IP destination address is considered a flow.

  - **TCP or UDP or SCTP Destination Port**—The class matches a single port or a contiguous range of ports. When you click **Next**, you are prompted to choose the protocol and enter the port number; click **...** to choose one already defined in ASDM.

Tip
For applications that use multiple, non-contiguous ports, use the Source and Destination IP Address (uses ACL) to match each port.

- **RTP Range**—The class map matches RTP traffic. When you click **Next**, you are prompted to enter an RTP port range, between 2000 and 65534. The maximum number of ports in the range is 16383.
- **IP DiffServ CodePoints (DSCP)**—The class matches up to eight DSCP values in the IP header. When you click **Next**, you are prompted to select or enter the desired values (move them into the Match on DSCP list).
- **IP Precedence**—The class map matches up to four precedence values, represented by the TOS byte in the IP header. When you click **Next**, you are prompted for the values.
- **Any Traffic**—Matches all traffic.

- **Add rule to existing traffic class**. If you already have a service policy rule on the same interface, or you are adding to the global service policy, this option lets you add an ACE to an existing ACL. You can add an ACE to any ACL that you previously created when you chose the Source and Destination IP Address (uses ACL) option for a service policy rule on this interface. For this traffic class, you can have only one set of rule actions even if you add multiple ACEs. You can add multiple ACEs to the same traffic class by repeating this entire procedure. When you click **Next**, you are prompted for the attributes of the access control entry.
- **Use an existing traffic class.** If you created a traffic class used by a rule on a different interface, you can reuse the traffic class definition for this rule. Note that if you alter the traffic class for one rule, the change is inherited by all rules that use that traffic class. If your configuration includes any **class-map** commands that you entered at the CLI, those traffic class names are also available (although to view the definition of the traffic class, you need to create the rule).
- **Use class default as the traffic class**. This option uses the class-default class, which matches all traffic. The class-default class is created automatically by the ASA and placed at the end of the policy. If you do not apply any actions to it, it is still created by the ASA, but for internal purposes only. You can apply actions to this class, if desired, which might be more convenient than creating a new traffic class that matches all traffic. You can only create one rule for this service policy using the class-default class, because each traffic class can only be associated with a single rule per service policy.

Step 4 If you selected a traffic matching criteria that requires additional configuration, enter the desired parameters and click **Next**.

Step 5 On the Rule Actions page, configure one or more rule actions. See Features Configured with Service Policies, on page 285 for a list of features and actions that you can apply, with pointers to additional details.

Step 6 Click **Finish**.

Add a Service Policy Rule for Management Traffic

To add a service policy rule for traffic directed to the ASA for management purposes, use the Add Service Policy Rule wizard. You will be asked to choose the scope of the policy, for a specific interface or global:

- Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with RADIUS accounting inspection, and an interface policy with connection limits, then both RADIUS accounting and connection limits are applied to the interface. However, if you

have a global policy with RADIUS accounting, and an interface policy with RADIUS accounting, then only the interface policy RADIUS accounting is applied to that interface.

- Global service policies provide default services to all interfaces. Unless overridden by an interface-specific policy, the global services are applied. By default, a global policy exists that includes a service policy rule for default application inspection. You can add a rule to the global policy using the wizard.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy Rules**, and click **Add** or **Add** > **Add Management Service Policy Rule**.

Step 2 In the Create a Service Policy and Apply To area:

a) Choose whether the policy applies to a specific **Interface** or **Global** to all interfaces.
b) If you select Interface, choose the name of the interface. If the interface already has a policy, then you are adding a rule to the existing policy.
c) If the interface does not already have a service policy, enter the name of the new policy.
d) (Optional) Enter a description for the policy.
e) Click **Next**.

Step 3 On the Traffic Classification Criteria page, choose one of the following options to specify the traffic to which to apply the policy actions and click **Next**.

- **Create a new traffic class**. Enter a traffic class name and an optional description.

  Identify the traffic using one of several criteria:

  - **Source and Destination IP Address (uses ACL)**—The class matches traffic specified by an extended ACL. When you click **Next**, you are prompted for the attributes of the access control entry, and the wizard builds the ACL. Optionally, you can select an existing ACL.

    When defining the ACE, the Match option creates a rule where traffic matching the addresses have actions applied. The Do Not Match option exempts the traffic from having the specified actions applied. For example, you want to match all traffic in 10.1.1.0/24 and apply connection limits to it, except for 10.1.1.25. In this case, create two rules, one for 10.1.1.0/24 using the Match option and one for 10.1.1.25 using the Do Not Match option. Be sure to arrange the rules so that the Do Not Match rule is above the Match rule, or else 10.1.1.25 will match the Match rule first.

  - **TCP or UDP or SCTP Destination Port**—The class matches a single port or a contiguous range of ports. When you click **Next**, you are prompted to choose the protocol and enter the port number; click **...** to choose one already defined in ASDM.

    Tip
    For applications that use multiple, non-contiguous ports, use the Source and Destination IP Address (uses ACL) to match each port.

- **Add rule to existing traffic class**. If you already have a service policy rule on the same interface, or you are adding to the global service policy, this option lets you add an ACE to an existing ACL. You can add an ACE to any ACL that you previously created when you chose the Source and Destination IP Address (uses ACL) option for a service policy rule on this interface. For this traffic class, you can have only one set of rule actions even if you add multiple ACEs. You can add multiple ACEs to the same

traffic class by repeating this entire procedure. When you click **Next**, you are prompted for the attributes of the access control entry.

- **Use an existing traffic class.** If you created a traffic class used by a rule on a different interface, you can reuse the traffic class definition for this rule. Note that if you alter the traffic class for one rule, the change is inherited by all rules that use that traffic class. If your configuration includes any **class-map** commands that you entered at the CLI, those traffic class names are also available (although to view the definition of the traffic class, you need to create the rule).

Step 4 If you selected a traffic matching criteria that requires additional configuration, enter the desired parameters and click **Next**.

Step 5 On the Rule Actions page, configure one or more rule actions.

- To configure RADIUS accounting inspection, choose an inspect map from the RADIUS Accounting Map drop-down list, or click **Configure** to add a map. See Features Configured with Service Policies, on page 285 for more information.
- To configure connection settings, see Configure Connection Settings for Specific Traffic Classes (All Services), on page 428.

Step 6 Click **Finish**.

Manage the Order of Service Policy Rules

The order of service policy rules on an interface or in the global policy affects how actions are applied to traffic. See the following guidelines for how a packet matches rules in a service policy:

- A packet can match only one rule in a service policy for each feature type.
- When the packet matches a rule that includes actions for a feature type, the ASA does not attempt to match it to any subsequent rules including that feature type.
- If the packet matches a subsequent rule for a different feature type, however, then the ASA also applies the actions for the subsequent rule.

For example, if a packet matches a rule for connection limits, and also matches a rule for application inspection, then both rule actions are applied.

If a packet matches a rule for application inspection, but also matches another rule that includes application inspection, then the second rule actions are not applied.

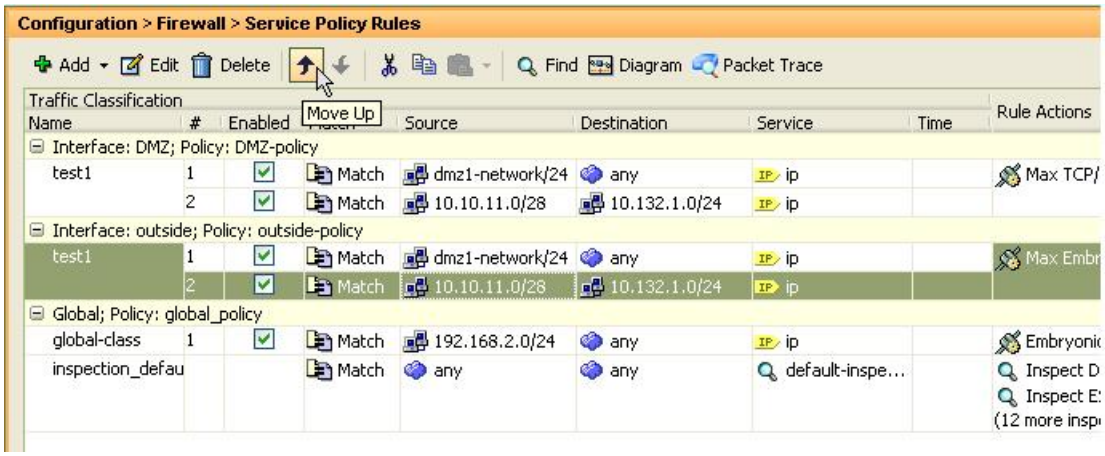
If your rule includes an ACL with multiple ACEs, then the order of ACEs also affects the packet flow. The ASA tests the packet against each ACE in the order in which the entries are listed. After a match is found, no more ACEs are checked. For example, if you create an ACE at the beginning of an ACL that explicitly permits all traffic, no further statements are ever checked.

To change the order of rules or ACEs within a rule, perform the following steps:

Procedure

Step 1 On the **Configuration** > **Firewall** > **Service Policy Rules** pane, choose the rule or ACE that you want to move up or down.

Step 2 Click the **Move Up** or **Move Down** button.



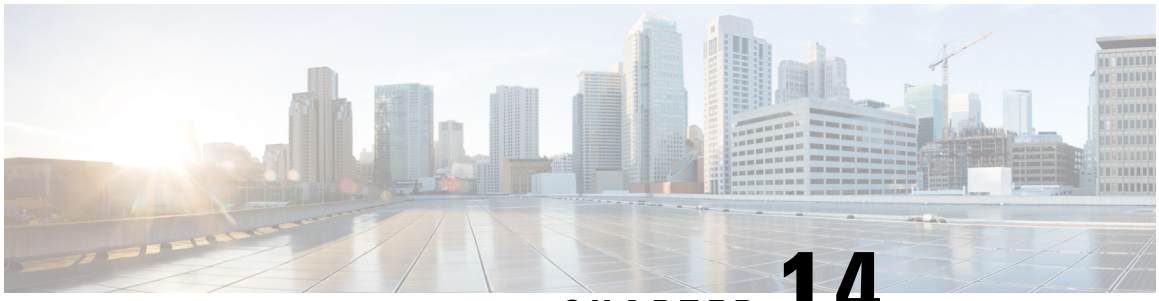
Note
If you rearrange ACEs in an ACL that is used in multiple service policies, then the change is inherited in all service policies.

Step 3 When you are done rearranging your rules or ACEs, click **Apply**.

History for Service Policies

Feature Name	Releases	Description
Modular Policy Framework	7.0(1)	Modular Policy Framework was introduced.
Management class map for use with RADIUS accounting traffic	7.2(1)	The management class map was introduced for use with RADIUS accounting traffic. The following commands were introduced: **class-map type management**, and **inspect radius-accounting**.
Inspection policy maps	7.2(1)	The inspection policy map was introduced. The following command was introduced: **class-map type inspect**.
Regular expressions and policy maps	7.2(1)	Regular expressions and policy maps were introduced to be used under inspection policy maps. The following commands were introduced: **class-map type regex**, **regex**, **match regex**.

Feature Name	Releases	Description
Match any for inspection policy maps	8.0(2)	The **match any** keyword was introduced for use with inspection policy maps: traffic can match one or more criteria to match the class map. Formerly, only **match all** was available.



CHAPTER 14

Getting Started with Application Layer Protocol Inspection

The following topics describe how to configure application layer protocol inspection.

- Application Layer Protocol Inspection, on page 299
- Configure Application Layer Protocol Inspection, on page 306
- Configure Regular Expressions, on page 310
- Monitoring Inspection Policies, on page 314
- History for Application Inspection, on page 315

Application Layer Protocol Inspection

Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection instead of passing the packet through the fast path. As a result, inspection engines can affect overall throughput. Several common inspection engines are enabled on the ASA by default, but you might need to enable others depending on your network.

The following topics explain application inspection in more detail.

When to Use Application Protocol Inspection

When a user establishes a connection, the ASA checks the packet against ACLs, creates an address translation, and creates an entry for the session in the fast path, so that further packets can bypass time-consuming checks. However, the fast path relies on predictable port numbers and does not perform address translations inside a packet.

Many protocols open secondary TCP or UDP ports. The initial session on a well-known port is used to negotiate dynamically assigned port numbers.

Other applications embed an IP address in the packet that needs to match the source address that is normally translated when it goes through the ASA.

If you use applications like these, then you need to enable application inspection.

When you enable application inspection for a service that embeds IP addresses, the ASA translates embedded addresses and updates any checksum or other fields that are affected by the translation.

When you enable application inspection for a service that uses dynamically assigned ports, the ASA monitors sessions to identify the dynamic port assignments, and permits data exchange on these ports for the duration of the specific session.

Inspection Policy Maps

You can configure special actions for many application inspections using an *inspection policy map*. These maps are optional: you can enable inspection for a protocol that supports inspection policy maps without configuring a map. These maps are needed only if you want something other than the default inspection actions.

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 criteria—You match application traffic to criteria specific to the application, such as a URL string, for which you then enable actions.

  For some traffic matching criteria, you use 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—Some inspection policy maps let you use an inspection class map to include multiple traffic matching criteria. You then identify the inspection class map in the inspection policy map and enable actions for the class as a whole. 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. However, you cannot set different actions for different matches.

- Parameters—Parameters affect the behavior of the inspection engine.

The following topics provide more details.

Replacing an In-Use Inspection Policy Map

If you have an inspection enabled with a policy map in a service policy, replacing the policy map is a two-step process. First, you must remove the inspection from the service policy and apply changes. Then, you add it back, select the new policy map name, and again apply changes.

How Multiple Traffic Classes are Handled

You can specify multiple inspection class maps or direct matches in the inspection policy map.

If a packet matches multiple different classes or direct matches, then the order in which the ASA applies the actions is determined by internal ASA rules, and not by the order they are added to the inspection 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.

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 criteria. If the first action is to log the packet, then a second action, such as resetting the connection, can occur.

If a packet matches multiple match criteria that are the same, then they are matched in the order they appear in the policy map.

A class map is determined to be the same type as another class map or direct match based on the lowest priority match option in the class map (the priority is based on the internal rules). If a class map has the same type of lowest priority match option 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 match for each class map is different, then the class map with the higher priority match option is matched first.

Guidelines for Application Inspection

Failover

State information for multimedia sessions that require inspection are not passed over the state link for stateful failover. The exceptions are GTP, M3UA, and SIP, which are replicated over the state link. You must configure strict application server process (ASP) state checking in M3UA inspection to get stateful failover.

Clustering

The following inspections are not supported in clustering:

- CTIQBE
- H323, H225, and RAS
- IPsec passthrough
- MGCP
- MMP
- RTSP
- SCCP (Skinny)
- WAAS

IPv6

Supports IPv6 for the following inspections:

- Diameter
- DNS over UDP
- FTP
- GTP
- HTTP
- ICMP
- IPsec pass-through
- IPv6
- M3UA
- SCCP (Skinny)

- SCTP
- SIP
- SMTP
- VXLAN

Supports NAT64 for the following inspections:

- DNS over UDP
- FTP
- HTTP
- ICMP
- SCTP

Additional Guidelines

- Some inspection engines do not support PAT, NAT, outside NAT, or NAT between same security interfaces. For more information about NAT support, see Default Inspections and NAT Limitations, on page 302.
- For all the application inspections, the ASA limits the number of simultaneous, active data connections to 200 connections. For example, if an FTP client opens multiple secondary connections, the FTP inspection engine allows only 200 active connections and the 201 connection is dropped and the adaptive security appliance generates a system error message.
- Inspected protocols are subject to advanced TCP-state tracking, and the TCP state of these connections is not automatically replicated. While these connections are replicated to the standby unit, there is a best-effort attempt to re-establish a TCP state.
- If the system determines that a TCP connection requires inspection, the system clears all TCP options except for the MSS and selective-acknowledgment (SACK) options on the packets before inspecting them. Other options are cleared even if you allow them in a TCP map applied to the connections.
- TCP/UDP Traffic directed to the ASA (to an interface) is inspected by default. However, ICMP traffic directed to an interface is never inspected, even if you enable ICMP inspection. Thus, a ping (echo request) to an interface can fail under specific circumstances, such as when the echo request comes from a source that the ASA can reach through a backup default route.

Defaults for Application Inspection

The following topics explain the default operations for application inspection.

Default Inspections and NAT Limitations

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.

The following table 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. In this table:

- Inspection engines that are enabled by default for the default port are in bold.
- The ASA is in compliance with the indicated standards, but it does not enforce compliance on packets being inspected. For example, FTP commands are supposed to be in a particular order, but the ASA does not enforce the order.

Table 12: Supported Application Inspection Engines

Application	Default Protocol, Port	NAT Limitations	Standards	Comments
CTIQBE	TCP/2748	No extended PAT. No NAT64. (Clustering) No static PAT.	—	—
DCERPC	TCP/135	No NAT64.	—	—
Diameter	TCP/3868 TCP/5868 (for TCP/TLS) SCTP/3868	No NAT/PAT.	RFC 6733	Requires the Carrier license.
DNS over UDP DNS over TCP	UDP/53 UDP/443 TCP/53	No NAT support is available for name resolution through WINS.	RFC 1123	You must enable DNS/TCP inspection in the DNS inspection policy map to inspect DNS over TCP. UDP/443 is used for Cisco Umbrella DNScrypt sessions only.
FTP	TCP/21	(Clustering) No static PAT.	RFC 959	—
GTP	UDP/3386 (GTPv0) UDP/2123 (GTPv1+)	No extended PAT. No NAT.	—	Requires the Carrier license.
H.323 H.225 and **RAS**	TCP/1720 UDP/1718 UDP (RAS) 1718-1719	(Clustering) No static PAT. No extended PAT. No NAT on same security interfaces. No NAT64.	ITU-T H.323, H.245, H225.0, Q.931, Q.932	—

Application	Default Protocol, Port	NAT Limitations	Standards	Comments
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	ICMP	—	—	ICMP traffic directed to an ASA interface is never inspected.
ICMP ERROR	ICMP	—	—	—
ILS (LDAP)	TCP/389	No extended PAT. No NAT64.	—	—
Instant Messaging (IM)	Varies by client	No extended PAT. No NAT64.	RFC 3860	—
IP Options	RSVP	No NAT64.	RFC 791, RFC 2113	—
IPsec Pass Through	UDP/500	No PAT. No NAT64.	—	—
IPv6	—	No NAT64.	RFC 2460	—
LISP	—	No NAT or PAT.	—	—
M3UA	SCTP/2905	No NAT or PAT for embedded addresses.	RFC 4666	Requires the Carrier license.
MGCP	UDP/2427, 2727	No extended PAT. No NAT64. (Clustering) No static PAT.	RFC 2705bis-05	—
MMP	TCP/5443	No extended PAT. No NAT64.	—	—
NetBIOS Name Server over IP	UDP/137, 138 (Source ports)	No extended PAT. No NAT64.	—	NetBIOS is supported by performing NAT of the packets for NBNS UDP port 137 and NBDS UDP port 138.
PPTP	TCP/1723	No NAT64. (Clustering) No static PAT.	RFC 2637	—
RADIUS Accounting	UDP/1646	No NAT64.	RFC 2865	—

Application	Default Protocol, Port	NAT Limitations	Standards	Comments
RSH	TCP/514	No PAT. No NAT64. (Clustering) No static PAT.	Berkeley UNIX	—
RTSP	TCP/554	No extended PAT. No NAT64. (Clustering) No static PAT.	RFC 2326, 2327, 1889	No handling for HTTP cloaking.
SCTP	SCTP	—	RFC 4960	Requires the Carrier license. Although you can do static network object NAT on SCTP traffic (no dyamic NAT/PAT), the inspection engine is not used for NAT.
SIP	TCP/5060 UDP/5060	No NAT/PAT on interfaces with the same, or lower to higher, security levels. No extended PAT. No NAT64 or NAT46. (Clustering) No static PAT.	RFC 2543	Does not handle TFTP uploaded Cisco IP Phone configurations under certain circumstances.
SKINNY (SCCP)	TCP/2000	No NAT on same security interfaces. No extended PAT. No NAT64, NAT46, or NAT66. (Clustering) No static PAT.	—	Does not handle TFTP uploaded Cisco IP Phone configurations under certain circumstances.
SMTP and **ESMTP**	TCP/25	No NAT64.	RFC 821, 1123	—
SNMP	UDP/161, 162 UDP/4161 on platforms that also run Secure Firewall eXtensible Operating System (FXOS).	No NAT or PAT.	RFC 1155, 1157, 1212, 1213, 1215	v.2 RFC 1902-1908; v.3 RFC 2570-2580. On FXOS platforms, if you configure SNMP, this inspection is enabled automatically, and you cannot disable it.
SQL*Net	TCP/1521	No extended PAT. No NAT64. (Clustering) No static PAT.	—	v.1 and v.2.

Application	Default Protocol, Port	NAT Limitations	Standards	Comments
STUN	TCP/3478 UDP/3478	(WebRTC) Static NAT/PAT44 only. (Cisco Spark) Static NAT/PAT44 and 64; and dynamic NAT/PAT.	RFC 5245, 5389	—
Sun RPC	TCP/111 UDP/111	No extended PAT. No NAT64.	—	—
TFTP	UDP/69	No NAT64. (Clustering) No static PAT.	RFC 1350	Payload IP addresses are not translated.
WAAS	TCP/1- 65535	No extended PAT. No NAT64.	—	—
XDMCP	UDP/177	No extended PAT. No NAT64. (Clustering) No static PAT.	—	—
VXLAN	UDP/4789	Not applicable	RFC 7348	Virtual Extensible Local Area Network.

Default Inspection Policy Maps

Some inspection types use hidden default policy maps. For example, if you enable ESMTP inspection without specifying a map, _default_esmtp_map is used.

The default inspection is described in the sections that explain each inspection type. You can view these default maps using the **show running-config all policy-map** command; use **Tools > Command Line Interface**.

DNS inspection is the only one that uses an explicitly-configured default map, preset_dns_map.

Configure Application Layer Protocol Inspection

You configure application inspection in service policies.

Inspection is enabled by default globally on all interfaces for some applications on their standard ports and protocols. See Default Inspections and NAT Limitations, on page 302 for more information on default inspections. A common method for customizing the inspection configuration is to customize the default global policy. You can alternatively create a new service policy as desired, for example, an interface-specific policy.

Before you begin

For some applications, you can perform special actions when you enable inspection by configuring inspection policy maps. The table later in this procedure shows which protocols allow inspection policy maps, with pointers to the instructions on configuring them. If you want to configure these advanced features, create the map before configuring inspection.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy Rules**.

Step 2 Open a rule.

- To edit the default global policy, select the "inspection_default" rule in the Global folder and click **Edit**.
- To create a new rule, click **Add > Add Service Policy Rule**. Proceed through the wizard to the Rules page.
- If you have another inspection rule, or a rule to which you are adding an inspection, select it and click **Edit**.

If you want to match non-standard ports, then create a new rule for the non-standard ports. See Default Inspections and NAT Limitations, on page 302 for the standard ports for each inspection engine.

You can combine multiple rules in the same service policy if desired, so you can create one rule to match certain traffic, and another to match different traffic. However, if traffic matches a rule that contains an inspection action, and then matches another rule that also has an inspection action, only the first matching rule is used.

If you are implementing RADIUS accounting inspection, create a management service policy rule instead. See Configure RADIUS Accounting Inspection, on page 400.

Step 3 On the Rule Actions wizard page or tab, select the **Protocol Inspection** tab.

Step 4 (To change an in-use policy) If you are editing any in-use policy to use a different inspection policy map, you must disable the inspection, and then re-enable it with the new inspection policy map name:

a) Uncheck the protocol's check box.
b) Click **OK**.
c) Click **Apply**.
d) Repeat these steps to return to the **Protocol Inspections** tab.

Step 5 Select the inspection type that you want to apply.

You can select multiple options on the default inspection traffic class only.

Some inspection engines let you control additional parameters when you apply the inspection to the traffic. Click **Configure** for the inspection type to configure an inspection policy map and other options. You can either choose an existing map, or create a new one. You can predefine inspection policy maps from the **Configuration > Firewall > Objects > Inspect Maps** list.

The following table lists the protocols you can inspect, whether they allow inspection policy maps or inspection class maps, and a pointer to detailed information about the inspection.

Table 13: Inspection Protocols

Protocol	Supports Inspection Policy Maps	Supports Inspection Class Maps	Notes
CTIQBE	No	No	See CTIQBE Inspection, on page 351.
DCERPC	Yes	Yes	See DCERPC Inspection, on page 317.

Protocol	Supports Inspection Policy Maps	Supports Inspection Class Maps	Notes
Diameter	Yes	Yes	See Diameter Inspection, on page 377. If you want to inspect encrypted Diameter traffic, choose **Enable encrypted traffic inspection** and select a TLS proxy (click **Manage** to create one if necessary).
DNS	Yes	Yes	See DNS Inspection, on page 320. If you are using the Botnet Traffic Filter, choose **Enable DNS snooping**. We suggest that you enable DNS snooping only on interfaces where external DNS requests are going. Enabling DNS snooping on all UDP DNS traffic, including that going to an internal DNS server, creates unnecessary load on the ASA. For example, if the DNS server is on the outside interface, you should enable DNS inspection with snooping for all UDP DNS traffic on the outside interface.
ESMTP	Yes	No	See SMTP and Extended SMTP Inspection, on page 341.
FTP	Yes	Yes	See FTP Inspection, on page 324. Select **Use Strict FTP** to select an inspection policy map. Strict FTP increases the security of protected networks by preventing web browsers from sending embedded commands in FTP requests.
GTP	Yes	No	See GTP Inspection Overview, on page 373.
H.323 H.225	Yes	Yes	See H.323 Inspection, on page 352.
H.323 RAS	Yes	Yes	See H.323 Inspection, on page 352.
HTTP	Yes	Yes	See HTTP Inspection, on page 328.
ICMP	No	No	See ICMP Inspection, on page 332.
ICMP Error	No	No	See ICMP Error Inspection, on page 332.
ILS	No	No	See ILS Inspection, on page 333.
IM	Yes	Yes	See Instant Messaging Inspection, on page 333.
IP-Options	Yes	No	See IP Options Inspection, on page 335.
IPSec Pass Thru	Yes	No	See IPsec Pass Through Inspection, on page 337.
IPv6	Yes	No	See IPv6 Inspection, on page 339.

Protocol	Supports Inspection Policy Maps	Supports Inspection Class Maps	Notes
LISP	Yes	No	For detailed information on configuring LISP, including inspection, see the clustering chapter in the general configuration guide.
M3UA	Yes	No	See M3UA Inspection, on page 378.
MGCP	Yes	No	See MGCP Inspection, on page 356.
NetBIOS	Yes	No	See NetBIOS Inspection, on page 340.
PPTP	No	No	See PPTP Inspection, on page 341.
RADIUS Accounting	Yes	No	See RADIUS Accounting Inspection Overview, on page 379. RADIUS accounting inspection is available for a management service policy only. You must select a policy map to implement this inspection.
RSH	No	No	See RSH Inspection, on page 341.
RTSP	Yes	No	See RTSP Inspection, on page 359.
SCCP (Skinny)	Yes	No	See Skinny (SCCP) Inspection, on page 366.
SCTP	Yes	No	See SCTP Application Layer Inspection, on page 376.
SIP	Yes	Yes	See SIP Inspection, on page 361. If you want to inspect encrypted SIP traffic, choose **Enable encrypted traffic inspection** and select a TLS proxy (click **Manage** to create one if necessary).
SNMP	Yes	No	See SNMP Inspection, on page 345.
SQLNET	No	No	See SQL*Net Inspection, on page 346.
STUN	No	No	See STUN Inspection, on page 369.
SUNRPC	No	No	See Sun RPC Inspection, on page 346. 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 SUNRPC inspection to that class.
TFTP	No	No	See TFTP Inspection, on page 348.
WAAS	No	No	Enables TCP option 33 parsing. Use when deploying Cisco Wide Area Application Services products.

Protocol	Supports Inspection Policy Maps	Supports Inspection Class Maps	Notes
XDMCP	No	No	See XDMCP Inspection, on page 348.
VXLAN	No	No	See VXLAN Inspection, on page 349.

Step 6 Click **OK** or **Finish** to save the service policy rule.

Configure Regular Expressions

Regular expressions define pattern matching for text strings. You can use these expressions in some protocol inspection maps to match packets based on strings such as URLs or the contents of particular header fields.

Create a Regular Expression

A regular expression matches text strings either literally as an exact string, or by using *metacharacters* so that you can match multiple variants of a text string. You can use a regular expression to match the content of certain application traffic; for example, you can match a URL string inside an HTTP packet.

Before you begin

See the **regex** command in the command reference for performance impact information when matching a regular expression to packets. In general, matching against long input strings, or trying to match a large number of regular expressions, will reduce system performance.

Note As an optimization, the ASA searches on the deobfuscated URL. Deobfuscation compresses multiple forward slashes (/) into a single slash. For strings that commonly use double slashes, like “http://”, be sure to search for “http:/” instead.

The following table lists the metacharacters that have special meanings.

Table 14: Regular Expression Metacharacters

Character	Description	Notes
.	Dot	Matches any single character. For example, **d.g** matches dog, dag, dtg, and any word that contains those characters, such as doggonnit.

Character	Description	Notes
(*exp*)	Subexpression	A subexpression segregates characters from surrounding characters, so that you can use other metacharacters on the subexpression. For example, **d(o|a)g** matches dog and dag, but **do|ag** matches do and ag. A subexpression can also be used with repeat quantifiers to differentiate the characters meant for repetition. For example, **ab(xy){3}z** matches abxyxyxyz.
|	Alternation	Matches either expression it separates. For example, **dog|cat** matches dog or cat.
?	Question mark	A quantifier that indicates that there are 0 or 1 of the previous expression. For example, **lo?se** matches lse or lose.
*	Asterisk	A quantifier that indicates that there are 0, 1 or any number of the previous expression. For example, **lo*se** matches lse, lose, loose, and so on.
+	Plus	A quantifier that indicates that there is at least 1 of the previous expression. For example, **lo+se** matches lose and loose, but not lse.
{*x*} or {*x*,}	Minimum repeat quantifier	Repeat at least *x* times. For example, **ab(xy){2,}z** matches abxyxyz, abxyxyxyz, and so on.
[*abc*]	Character class	Matches any character in the brackets. For example, **[abc]** matches a, b, or c.
[^*abc*]	Negated character class	Matches a single character that is not contained within the brackets. For example, **[^abc]** matches any character other than a, b, or c. **[^A-Z]** matches any single character that is not an uppercase letter.
[*a*-*c*]	Character range class	Matches any character in the range. **[a-z]** matches any lowercase letter. You can mix characters and ranges: **[abcq-z]** matches a, b, c, q, r, s, t, u, v, w, x, y, z, and so does [**a-cq-z]**. The dash (-) character is literal only if it is the last or the first character within the brackets: **[abc-]** or **[-abc]**.
“”	Quotation marks	Preserves trailing or leading spaces in the string. For example, **“ test”** preserves the leading space when it looks for a match.
^	Caret	Specifies the beginning of a line.
\	Escape character	When used with a metacharacter, matches a literal character. For example, **\[** matches the left square bracket.

Character	Description	Notes
char	Character	When character is not a metacharacter, matches the literal character.
\r	Carriage return	Matches a carriage return 0x0d.
\n	Newline	Matches a new line 0x0a.
\t	Tab	Matches a tab 0x09.
\f	Formfeed	Matches a form feed 0x0c.
\x*NN*	Escaped hexadecimal number	Matches an ASCII character using hexadecimal (exactly two digits).
NNN	Escaped octal number	Matches an ASCII character as octal (exactly three digits). For example, the character 040 represents a space.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Regular Expressions**.

Step 2 In the Regular Expressions area, do one of the following:

- Choose **Add** to add a new object. Enter a name and optionally, a description.
- Choose an existing object and click **Edit**.

Step 3 Either enter the regular expression in the **Value** field, or click **Build** to get help creating the expression.

The regular expression is limited to 100 characters in length.

If you click **Build**, use the following process to create the expression:

a) In the Build Snippet area, create a component of the expression using the following options. Look at the Snippet Preview area at the end of this section to see the expression you are building.

- Starts at the beginning of the line (^)—Indicates that the snippet should start at the beginning of a line, using the caret (^) metacharacter. Be sure to insert any snippet with this option at the beginning of the regular expression.
- Specify Character String—If you are trying to match a specific string, such as a word or phrase, enter the string.

  If there are any metacharacters in your text string that you want to be used literally, choose **Escape Special Characters** to add the backslash (\) escape character before them. for example, if you enter “example.com,” this option converts it to “example\.com”.

  If you want to match upper and lower case characters, choose **Ignore Case**. For example, “cats” is converted to “[cC][aA][tT][sS]”.
- Specify Character—If you are trying to match a specific type of character or set of characters, rather than a particular phrase, select this option and identify the characters using these options:

- **Negate the character**—Specifies not to match the character you identify.
- **Any character (.)**—Inserts the period (.) metacharacter to match any character. For example, **d.g** matches dog, dag, dtg, and any word that contains those characters, such as doggonnit.
- **Character set**—Inserts a character set. Text can match any character in the set. For example, if you specify [0-9A-Za-z], then this snippet will match any character from A to Z (upper or lower case) or any digit 0 through 9. The [\n\f\r\t] set matches a new line, form feed, carriage return, or a tab.
- **Special character**—Inserts a character that requires an escape, including \, ?, *, +, |, ., [, (, or ^. The escape character is the backslash (\), which is automatically entered when you choose this option.
- **Whitespace character**—Whitespace characters include \n (new line), \f (form feed), \r (carriage return), or \t (tab).
- **Three digit octal number**—Matches an ASCII character as octal (up to three digits). For example, the character \040 represents a space. The backslash (\) is entered automatically.
- **Two digit hexadecimal number**—Matches an ASCII character using hexadecimal (exactly two digits). The backslash (\) is entered automatically.
- **Specified character**—Enter any single character.

b) Add the snippet to the regular expression box using one of the following buttons. Note that you can also type directly in the regular expression.
   - **Append Snippet**—Adds the snippet to the end of the regular expression.
   - **Append Snippet as Alternate**—Adds the snippet to the end of the regular expression separated by a pipe (|), which matches either expression it separates. For example, **dog|cat** matches dog or cat.
   - **Insert Snippet at Cursor**—Inserts the snippet at the cursor.

c) Repeat the process to add snippets until the expression is complete.

d) (Optional.) In **Selection Occurrences**, select how often the expression or parts of it must match text to be considered a match. Select text in the Regular Expression field, click one of the following options, and then click **Apply to Selection**. For example, if the regular expression is “test me,” and you select “me” and apply **One or more times**, then the regular expression changes to “test (me)+”.
   - **Zero or one times (?)**—There are 0 or 1 of the previous expression. For example, **lo?se** matches lse or lose.
   - **One or more times (+)**—There is at least 1 of the previous expression. For example, **lo+se** matches lose and loose, but not lse.
   - **Any number of times (*)**—There are 0, 1 or any number of the previous expression. For example, **lo*se** matches lse, lose, loose, and so on.
   - **At least**—Repeat at least *x* times. For example, **ab(xy){2,}z** matches abxyxyz, abxyxyxyz, and so on.
   - **Exactly**—Repeat exactly *x* times. For example, **ab(xy){3}z** matches abxyxyxyz.

e) Click **Test** to verify your expression will match the intended text. If the test is unsuccessful, you can try editing it in the test dialog, or return to the expression builder. If you edit the expression in the text dialog and click **OK**, the edits are saved and reflected in the expression builder.
f) Click **OK**.

Create a Regular Expression Class Map

A regular expression class map identifies one or more regular expression. It is simply a collection of regular expression objects. You can use a regular expression class map in many cases in replace of a regular expression object.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Regular Expressions**.
Step 2 In the Regular Expressions Classes area, do one of the following:

- Choose **Add** to add a new class map. Enter a name and optionally, a description.
- Choose an existing class map and click **Edit**.

Step 3 Select the expressions you want in the map and click **Add**. Remove any you do not want.
Step 4 Click **OK**.

Monitoring Inspection Policies

To monitor inspection service policies, enter the following commands. Select **Tools** > **Command Line Interface** to enter these commands. See the command reference on Cisco.com for detailed syntax and examples.

- **show service-policy inspect** *protocol*

  Displays statistics for inspection service policies. The *protocol* is the protocol from the inspect command, for example **dns**. However, not all inspection protocols show statistics with this command. For example:

```
asa# show service-policy inspect dns

Global policy:
  Service-policy: global_policy
    Class-map: inspection_default
      Inspect: dns preset_dns_map, packet 0, lock fail 0, drop 0, reset-drop 0,
5-min-pkt-rate 0 pkts/sec, v6-fail-close 0
        message-length maximum client auto, drop 0
        message-length maximum 512, drop 0
        dns-guard, count 0
        protocol-enforcement, drop 0
        nat-rewrite, count 0
asa#
```

- **show conn**

  Shows current connections for traffic passing through the device. This command has a wide range of keywords so that you can get information about various protocols.

- Additional commands for specific inspected protocols:
  - **show ctiqbe**

    Displays information about the media connections allocated by the CTIQBE inspection engine

  - **show h225**

    Displays information for H.225 sessions.

  - **show h245**

    Displays information for H.245 sessions established by endpoints using slow start.

  - **show h323 ras**

    Displays connection information for H.323 RAS sessions established between a gatekeeper and its H.323 endpoint.

  - **show mgcp** {**commands** | **sessions** }

    Displays the number of MGCP commands in the command queue or the number of existing MGCP sessions.

  - **show sip**

    Displays information for SIP sessions.

  - **show skinny**

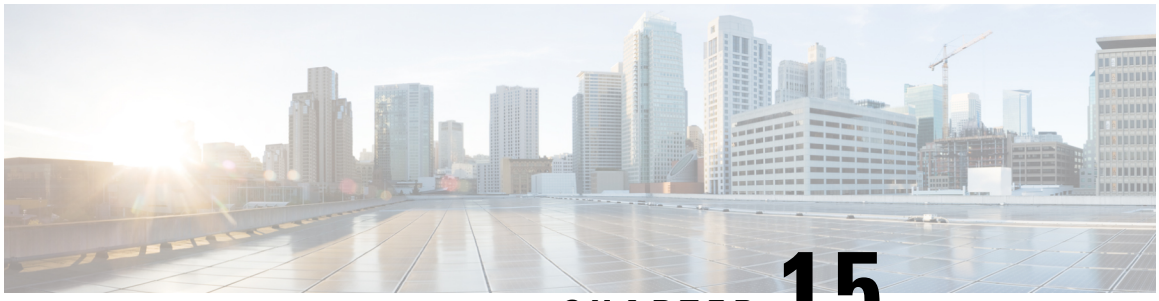
    Displays information for Skinny (SCCP) sessions.

  - **show sunrpc-server active**

    Displays the pinholes opened for Sun RPC services.

History for Application Inspection

Feature Name	Releases	Description
Inspection policy maps	7.2(1)	The inspection policy map was introduced. The following command was introduced: **class-map type inspect**.
Regular expressions and policy maps	7.2(1)	Regular expressions and policy maps were introduced to be used under inspection policy maps. The following commands were introduced: **class-map type regex**, **regex**, **match regex**.
Match any for inspection policy maps	8.0(2)	The **match any** keyword was introduced for use with inspection policy maps: traffic can match one or more criteria to match the class map. Formerly, only **match all** was available.



CHAPTER 15

Inspection of Basic Internet Protocols

The following topics explain application inspection for basic Internet protocols. For information on why you need to use inspection for certain protocols, and the overall methods for applying inspection, see Getting Started with Application Layer Protocol Inspection, on page 299.

- DCERPC Inspection, on page 317
- DNS Inspection, on page 320
- FTP Inspection, on page 324
- HTTP Inspection, on page 328
- ICMP Inspection, on page 332
- ICMP Error Inspection, on page 332
- ILS Inspection, on page 333
- Instant Messaging Inspection, on page 333
- IP Options Inspection, on page 335
- IPsec Pass Through Inspection, on page 337
- IPv6 Inspection, on page 339
- NetBIOS Inspection, on page 340
- PPTP Inspection, on page 341
- RSH Inspection, on page 341
- SMTP and Extended SMTP Inspection, on page 341
- SNMP Inspection, on page 345
- SQL*Net Inspection, on page 346
- Sun RPC Inspection, on page 346
- TFTP Inspection, on page 348
- XDMCP Inspection, on page 348
- VXLAN Inspection, on page 349
- History for Basic Internet Protocol Inspection, on page 349

DCERPC Inspection

DCERPC inspection is not enabled in the default inspection policy, so you must enable it if you need this inspection. You can simply edit the default global inspection policy to add DCERPC inspection. You can alternatively create a new service policy as desired, for example, an interface-specific policy.

The following sections describe the DCERPC inspection engine.

DCERPC Overview

Microsoft Remote Procedure Call (MSRPC), based on 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.

The DCERPC inspection engine inspects 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 configurable timeouts.

DCE inspection supports the following universally unique identifiers (UUIDs) and messages:

- End point mapper (EPM) UUID. All EPM messages are supported.
- ISystemMapper UUID (non-EPM). Supported messages are:
  - RemoteCreateInstance opnum4
  - RemoteGetClassObject opnum3
- OxidResolver UUID (non-EPM). Supported message is:
  - ServerAlive2 opnum5
- Any message that does not contain an IP address or port information because these messages do not require inspection.

Configure a DCERPC Inspection Policy Map

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.

When defining traffic matching criteria, you can either create a class map or include the match statements 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 you can reuse class maps. The following procedure covers inspection policy maps, but also explains the traffic matching criteria available in the class map. To create a class map, select **Configuration** > **Firewall** > **Objects** > **Class Maps** > **DCERPC**.



Tip You can configure inspection maps while creating service policies, in addition to the procedure explained below. The contents of the map are the same regardless of how you create it.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **DCERPC**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map to view its contents. You can change the security level directly, or click **Customize** to edit the map. The remainder of the procedure assumes you are customizing or adding a map.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 In the **Security Level** view of the DCERPC Inspect Map dialog box, select the level that best matches your desired configuration.

If one of the preset levels matches your requirements, you are now done. Just click **OK**, skip the rest of this procedure, and use the map in a service policy rule for DCERPC inspection.

If you need to customize the settings further, click **Details** and continue with the procedure.

Tip
The **UUID Filtering** button is a shortcut to configure message filtering, which is explained later in this procedure.

Step 5 Configure the desired options.

- **Pinhole Timeout**—Sets the pinhole timeout. Because a client may use the server information returned by the endpoint mapper for multiple connections, the timeout value is configurable based on the client application environment. Range is from 0:0:1 to 1193:0:0.
- **Enforce endpoint-mapper service**—Whether to enforce the endpoint mapper service during binding so that only its service traffic is processed.
- **Enable endpoint-mapper service lookup**—Whether to enable the lookup operation of the endpoint mapper service. You can also enforce a timeout for the service lookup. If you do not configure a timeout, the pinhole timeout is used.

Step 6 (Optional.) Click the **Inspections** tab and define the actions to take for specific types of messages.

You can define traffic matching criteria based on DCERPC class maps, by configuring matches directly in the inspection map, or both.

a) Do any of the following:
   - Click **Add** to add a new criterion.
   - Select an existing criterion and click **Edit**.

b) Choose **Single Match** to define the criterion directly, or **Multiple Match**, in which case you select the DCERPC class map that defines the criteria.

c) If you are defining the criterion here, choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). Then, select the desired UUID:
   - **ms-rpc-epm**—Matches Microsoft RPC EPM messages.
   - **ms-rpc-isystemactivator**—Matches ISystemMapper messages.
   - **ms-rpc-oxidresolver**—Matches OxidResolver messages.

d) Choose whether to **Reset** or **Log** the connection. You can also enable logging if you elect to reset the connection. Resetting the connection drops the packet, closes the connection, and sends a TCP reset to the server or client.
e) Click **OK** to add the criterion. Repeat the process as needed.

Step 7 Click **OK**.

You can now use the inspection map in a DCERPC inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

DNS Inspection

DNS inspection is enabled by default. You need to configure it only if you want non-default processing. The following sections describe DNS application inspection.

Defaults for DNS Inspection

DNS inspection is enabled by default, using the preset_dns_map inspection class map:

- The maximum DNS message length is 512 bytes.
- DNS over TCP inspection is disabled.
- The maximum client DNS message length is automatically set to match the Resource Record.
- DNS Guard is enabled, so the ASA tears down the DNS session associated with a DNS query as soon as the DNS reply is forwarded by the ASA. The ASA also monitors the message exchange to ensure that the ID of the DNS reply matches the ID of the DNS query.
- Translation of the DNS record based on the NAT configuration is enabled.
- Protocol enforcement is enabled, which enables DNS message format check, including domain name length of no more than 255 characters, label length of 63 characters, compression, and looped pointer check.

Note Elevated CPU utilization is likely to occur in Cisco ASA when a large number of packets is inspected per second. As the number of packets being processed increases—especially when the device nears its maximum throughput or connection capacity—the CPU usage too rises. This occurs because ASA must perform more operations per second, which increases CPU resource consumption.

Configure DNS Inspection Policy Map

You can create a DNS inspection policy map to customize DNS inspection actions if the default inspection behavior is not sufficient for your network.

You can optionally create a DNS inspection class map to define the traffic class for DNS inspection. The other option is to define the traffic classes directly in the DNS inspection policy map. The difference between creating a class map and defining the traffic match directly in the inspection map is that you can create more complex match criteria and you can reuse class maps. Although this procedure explains inspection maps, the matching criteria used in class maps are the same as those explained in the step relating to the **Inspection** tab. You can configure DNS class maps by selecting **Configuration** > **Firewall** > **Objects** > **Class Maps** > **DNS**, or by creating them while configuring the inspection map.

Tip You can configure inspection maps while creating service policies, in addition to the procedure explained below. The contents of the map are the same regardless of how you create it.

Before you begin

Some traffic matching options use regular expressions for matching purposes. If you intend to use one of those techniques, first create the regular expression or regular expression class map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **DNS**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map to view its contents. You can change the security level directly, or click **Customize** to edit the map. The remainder of the procedure assumes you are customizing or adding a map.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 In the **Security Level** view of the DNS Inspect Map dialog box, select the level that best matches your desired configuration. The default level is Low.

If one of the preset levels matches your requirements, you are now done. Just click **OK**, skip the rest of this procedure, and use the map in a service policy rule for DNS inspection.

If you need to customize the settings further, click **Details**, and continue with the procedure.

Step 5 Click the **Protocol Conformance** tab and choose the desired options:

- **Enable DNS guard function**—Using DNS Guard, the ASA tears down the DNS session associated with a DNS query as soon as the DNS reply is forwarded by the ASA. The ASA also monitors the message exchange to ensure that the ID of the DNS reply matches the ID of the DNS query.
- **Enable NAT re-write function**—Translates the DNS record based on the NAT configuration.

- **Enable protocol enforcement**—Enables DNS message format check, including domain name length of no more than 255 characters, label length of 63 characters, compression, and looped pointer check.
- **Randomize the DNS identifier for DNS query**.
- **Enable TCP inspection**—Enables inspection of DNS over TCP traffic. Ensure that DNS/TCP port 53 traffic is part of the class to which you apply DNS inspection. The inspection default class includes TCP/53.
- **Enforce TSIG resource record to be present in DNS message**—You can drop or log non-conforming packets, and optionally log dropped packets.

Step 6 Click the **Filtering** tab and choose the desired options.

- Global Settings—Choose whether to drop packets that exceed the specified maximum length regardless of whether they are from the client or server, from 512 to 65535 bytes.
- Server Settings—**Drop packets that exceed specified maximum length** and **Drop packets sent to server that exceed length indicated by the RR**—Sets the maximum server DNS message length, from 512 to 65535 bytes, or sets the maximum length to the value in the Resource Record. If you enable both settings, the lower value is used.
- Client Settings—**Drop packets that exceed specified maximum length** and **Drop packets sent to server that exceed length indicated by the RR**—Sets the maximum client DNS message length, from 512 to 65535 bytes, or sets the maximum length to the value in the Resource Record. If you enable both settings, the lower value is used.

Step 7 Click the **Mismatch Rate** tab and choose whether to enable logging when the DNS ID mismatch rate exceeds the specified threshold. For example, you could set a threshold of 30 mismatches per 3 seconds.

Step 8 Click the **Inspections** tab and define the specific inspections you want to implement based on traffic characteristics.

You can define traffic matching criteria based on DNS class maps, by configuring matches directly in the inspection map, or both.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose **Single Match** to define the criterion directly, or **Multiple Match**, in which case you select the DNS class map that defines the criteria.

c) If you are defining the criterion here, choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). For example, if No Match is selected on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map. Then, configure the criterion as follows:

- Header Flag—Select whether the flag should equal or contain the specified value, then either select the header flag name, or enter the hex value of the header (0x0 to 0xfff). If you select multiple header values, “equals” requires that all flags are present, “contains” that any one of the flags is present, in the packet. Header flag names are **AA** (Authoritative Answer), **QR** (Query), **RA** (Recursion Available), **RD** (Recursion Desired), **TC** (Truncation).

- Type—The DNS Type field name or value in the packet. Field names are **A** (IPv4 address), **AXFR** (full zone transfer), **CNAME** (canonical name), **IXFR** (incremental zone transfer), **NS** (authoritative name server), **SOA** (start of a zone of authority) or **TSIG** (transaction signature). Values are arbitrary numbers in the DNS Type field from 0 to 65535: either enter a specific value or a range of values.
- Class—The DNS Class field name or value in the packet. Internet is the only possible field name. Values are arbitrary numbers in the DNS Class field from 0 to 65535: either enter a specific value or a range of values.
- Question—The question portion of a DNS message.
- Resource Record—The DNS resource record. Choose whether to match the additional, answer, or authority resource record section.

d) Choose the primary action to take for matching traffic: drop packet, drop connection, mask (for Header Flag matches only) or none.
e) Choose whether to enable or disable logging. You must disable logging if you want to enforce TSIG.
f) Chose whether to enforce the presence of a TSIG resource record. You can drop the packet, log it, or drop and log it. Usually, you must select **Primary Action: None** and **Log: Disable** to enforce TSIG. However, for Header Flag matches, you can enforce TSIG along with the mask primary action.
g) Click **OK** to add the inspection. Repeat the process as needed.

Step 9 Click the **Umbrella Connections** tab and enable the connection to Cisco Umbrella in the cloud.

The options on this tab work only if you configure the Cisco Umbrella connection on the **Configuration** > **Firewall** > **Objects** > **Umbrella** page. You must then configure the options on this tab to get the device to register with Cisco Umbrella, so that the device can redirect DNS lookups to Cisco Umbrella. Cisco Umbrella can then apply your FQDN-based security policies. For more information, see Cisco Umbrella, on page 125.

- **Umbrella**—Enables Cisco Umbrella. You can optionally specify the name of the Cisco Umbrella policy to apply to the device in the **Umbrella Tag** field. If you do not specify a policy, the default policy is applied. After registration, the Umbrella device ID is displayed next to the tag.
- **Enable Dnscrypt**—Enables DNScrypt to encrypt connections between the device and Cisco Umbrella. Enabling DNScrypt starts the key-exchange thread with the Umbrella resolver. The key-exchange thread performs the handshake with the resolver every hour and updates the device with a new secret key. Because DNScrypt uses UDP/443, you must ensure that the class map used for DNS inspection includes that port. Note that the default inspection class already includes UDP/443 for DNS inspection.
- **Fail Open**—Enable fail open if you want DNS resolution to work if the Umbrella DNS server is unavailable. When failing open, if the Cisco Umbrella DNS server is unavailable, Umbrella disables itself on this policy map and allows DNS requests to go to the other DNS servers configured on the system, if any. When the Umbrella DNS servers are available again, the policy map resumes using them. If you do not select this option, DNS requests continue to go to the unreachable Umbrella resolver, so they will not get a response

Step 10 Click **OK** in the DNS Inspect Map dialog box.

You can now use the inspection map in a DNS inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

FTP Inspection

FTP inspection is enabled by default. You need to configure it only if you want non-default processing. The following sections describe the FTP inspection engine.

FTP Inspection Overview

The FTP application inspection inspects the FTP sessions and performs four tasks:

- Prepares dynamic secondary data connection 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.
- Tracks the FTP command-response sequence.
- Generates an audit trail.
  - Audit record 303002 is generated for each file that is retrieved or uploaded.
  - Audit record 201005 is generated if the secondary dynamic channel preparation failed due to memory shortage.
- Translates the embedded IP address.



Note If you disable FTP inspection, outbound users can start connections only in passive mode, and all inbound FTP is disabled.

Strict FTP

Strict FTP increases the security of protected networks by preventing web browsers from sending embedded commands in FTP requests. To enable strict FTP, click the **Configure** button next to FTP on the Configuration > Firewall > Service Policy Rules > Edit Service Policy Rule > Rule Actions > Protocol Inspection tab.

When you use strict FTP, you can optionally specify an FTP inspection policy map to specify FTP commands that are not permitted to pass through the ASA.

Strict FTP inspection enforces the following behavior:

- An FTP command must be acknowledged before the ASA allows a new command.
- The ASA drops connections that send embedded commands.
- The 227 and PORT commands are checked to ensure they do not appear in an error string.

Caution Using strict FTP may cause the failure of FTP clients that are not strictly compliant with FTP RFCs. Additionally, you must ensure you apply the inspection to your FTP ports only (TCP/21 is the normal FTP port). Strict FTP inspection applied to non-FTP traffic can result in unexpected traffic loss, especially HTTP traffic.

With strict FTP inspection, each FTP command and response sequence is tracked for the following anomalous activity:

- Truncated command—Number of commas in the PORT and PASV reply command is checked to see if it is five. If it is not five, then the PORT command is assumed to be truncated and the TCP connection is closed.
- Incorrect command—Checks the FTP command to see if it ends with <CR><LF> characters, as required by the RFC. If it does not, the connection is closed.
- Size of RETR and STOR commands—These are checked against a fixed constant. If the size is greater, then an error message is logged and the connection is closed.
- Command spoofing—The PORT command should always be sent from the client. The TCP connection is denied if a PORT command is sent from the server.
- Reply spoofing—PASV reply command (227) should always be sent from the server. The TCP connection is denied if a PASV reply command is sent from the client. This prevents the security hole when the user executes “227 xxxxx a1, a2, a3, a4, p1, p2.”
- TCP stream editing—The ASA closes the connection if it detects TCP stream editing.
- Invalid port negotiation—The negotiated dynamic port value is checked to see if it is less than 1024. As port numbers in the range from 1 to 1024 are reserved for well-known connections, if the negotiated port falls in this range, then the TCP connection is freed.
- Command pipelining—The number of characters present after the port numbers in the PORT and PASV reply command is cross checked with a constant value of 8. If it is more than 8, then the TCP connection is closed.
- The ASA replaces the FTP server response to the SYST command with a series of Xs to prevent the server from revealing its system type to FTP clients. To override this default behavior, use the **no mask-syst-reply** command in the FTP map.

Configure an FTP Inspection Policy Map

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 inspection policy map. You can then apply the map when you enable FTP inspection.

You can optionally create an FTP inspection class map to define the traffic class for FTP inspection. The other option is to define the traffic classes directly in the FTP inspection policy map. The difference between creating a class map and defining the traffic match directly in the inspection map is that you can create more complex match criteria and you can reuse class maps. Although this procedure explains inspection maps, the matching criteria used in class maps are the same as those explained in the step relating to the **Inspection** tab. You can configure DNS class maps by selecting **Configuration** > **Firewall** > **Objects** > **Class Maps** > **FTP**, or by creating them while configuring the inspection map.

Tip You can configure inspection maps while creating service policies, in addition to the procedure explained below. The contents of the map are the same regardless of how you create it.

Before you begin

Some traffic matching options use regular expressions for matching purposes. If you intend to use one of those techniques, first create the regular expression or regular expression class map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **FTP**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map to view its contents. You can change the security level directly, or click **Customize** to edit the map. The remainder of the procedure assumes you are customizing or adding a map.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 In the **Security Level** view of the FTP Inspect Map dialog box, select the level that best matches your desired configuration. The default level is High.

If one of the preset levels matches your requirements, you are now done. Just click **OK**, skip the rest of this procedure, and use the map in a service policy rule for FTP inspection.

If you need to customize the settings further, click **Details**, and continue with the procedure.

Tip
The **File Type Filtering** button is a shortcut to configure file media or MIME type inspection, which is explained later in this procedure.

Step 5 Click the **Parameters** tab and choose whether to mask the greeting banner from the server or mask the reply to the SYST command.

Masking these items prevents the client from discovering server information that might be helpful in an attack.

Step 6 Click the **Inspections** tab and define the specific inspections you want to implement based on traffic characteristics.

You can define traffic matching criteria based on FTP class maps, by configuring matches directly in the inspection map, or both.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose **Single Match** to define the criterion directly, or **Multiple Match**, in which case you select the FTP class map that defines the criteria.

c) If you are defining the criterion here, choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map. Then, configure the criterion as follows:

- File Name—Match the name of the file being transferred against the selected regular expression or regular expression class.
- File Type—Match the MIME or media type of the file being transferred against the selected regular expression or regular expression class.
- Server—Match the FTP server name against the selected regular expression or regular expression class.
- User—Match the name of the logged-in user against the selected regular expression or regular expression class.
- Request Command—The FTP command used in the packet, any combination of the following:
  - **APPE**—Append to a file.
  - **CDUP**—Changes to the parent directory of the current working directory.
  - **DELE**—Delete a file on the server.
  - **GET**—Gets a file from the server.
  - **HELP**—Provides help information.
  - **MKD**—Makes a directory on the server.
  - **PUT**—Sends a file to the server.
  - **RMD**—Deletes a directory on the server.
  - **RNFR**—Specifies the "rename-from" filename.
  - **RNTO**—Specifies the "rename-to" filename.
  - **SITE**—Used to specify a server-specific command. This is usually used for remote administration.
  - **STOU**—Stores a file using a unique file name.

d) Choose whether to enable or disable logging. The action is always to reset the connection, which drops the packet, closes the connection, and sends a TCP reset to the server or client.

e) Click **OK** to add the inspection. Repeat the process as needed.

Step 7 Click **OK** in the FTP Inspect Map dialog box.

You can now use the inspection map in a FTP inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

HTTP Inspection

If you are not using a purpose-built module for HTTP inspection and application filtering, such as ASA FirePOWER, you can manually configure HTTP inspection on the ASA.

HTTP inspection is not enabled in the default inspection policy, so you must enable it if you need this inspection. However, the default inspect class does include the default HTTP ports, so you can simply edit the default global inspection policy to add HTTP inspection. You can alternatively create a new service policy as desired, for example, an interface-specific policy.

Tip Do not configure HTTP inspection in both a service module and on the ASA, as the inspections are not compatible.

The following sections describe the HTTP inspection engine.

HTTP Inspection Overview

Tip You can install a service module that performs application and URL filtering, which includes HTTP inspection, such as ASA FirePOWER. The HTTP inspection running on the ASA is not compatible with these modules. Note that it is far easier to configure application filtering using a purpose-built module rather than trying to manually configure it on the ASA using an HTTP inspection policy map.

Use the HTTP inspection engine to protect against specific attacks and other threats that are associated with HTTP traffic.

HTTP application inspection scans HTTP headers and body, and performs various checks on the data. These checks prevent various HTTP constructs, content types, and tunneling and messaging protocols from traversing the security appliance.

The enhanced HTTP inspection feature, which is also known as an application firewall and is available when you configure an HTTP inspection policy map, can help prevent attackers from using HTTP messages for circumventing network security policy.

HTTP application inspection can block tunneled applications and non-ASCII characters in HTTP requests and responses, preventing malicious content from reaching the web server. Size limiting of various elements in HTTP request and response headers, URL blocking, and HTTP server header type spoofing are also supported.

Enhanced HTTP inspection verifies the following for all HTTP messages:

- Conformance to RFC 2616
- Use of RFC-defined methods only.
- Compliance with the additional criteria.

Configure an HTTP Inspection Policy Map

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.

You can optionally create an HTTP inspection class map to define the traffic class for HTTP inspection. The other option is to define the traffic classes directly in the HTTP inspection policy map. The difference between creating a class map and defining the traffic match directly in the inspection map is that you can create more complex match criteria and you can reuse class maps. Although this procedure explains inspection maps, the matching criteria used in class maps are the same as those explained in the step relating to the **Inspection** tab. You can configure HTTP class maps by selecting **Configuration** > **Firewall** > **Objects** > **Class Maps** > **HTTP**, or by creating them while configuring the inspection map.

Tip You can configure inspection maps while creating service policies, in addition to the procedure explained below. The contents of the map are the same regardless of how you create it.

Before you begin

Some traffic matching options use regular expressions for matching purposes. If you intend to use one of those techniques, first create the regular expression or regular expression class map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **HTTP**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map to view its contents. You can change the security level directly, or click **Customize** to edit the map. The remainder of the procedure assumes you are customizing or adding a map.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 In the **Security Level** view of the HTTP Inspect Map dialog box, select the level that best matches your desired configuration. The default level is Low.

If one of the preset levels matches your requirements, you are now done. Just click **OK**, skip the rest of this procedure, and use the map in a service policy rule for HTTP inspection.

If you need to customize the settings further, click **Details**, and continue with the procedure.

Tip
The **URI Filtering** button is a shortcut to configure Request URI inspection, which is explained later in this procedure.

Step 5 Click the **Parameters** tab and configure the desired options.

- **Body Match Maximum**—The maximum number of characters in the body of an HTTP message that should be searched in a body match. Default is 200 bytes. A large number will have a significant impact on performance.
- **Check for protocol violations**—Whether to verify that packets conform to the HTTP protocol. For violations, you can drop the connection, reset it, or log it. When dropping or resetting, you can also enable logging.
- **Spoof server string**—Replaces the server HTTP header value with the specified string, up to 82 characters.

Step 6 Click the **Inspections** tab and define the specific inspections you want to implement based on traffic characteristics.

You can define traffic matching criteria based on HTTP class maps, by configuring matches directly in the inspection map, or both.

a) Do any of the following:
   - Click **Add** to add a new criterion.
   - Select an existing criterion and click **Edit**.

b) Choose **Single Match** to define the criterion directly, or **Multiple Match**, in which case you select the HTTP class map that defines the criteria.

c) If you are defining the criterion here, choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map. Then, configure the criterion as follows:
   - Request/Response Content Type Mismatch—Match packets where the content type in the response does not match one of the MIME types in the accept field of the request.
   - Request Arguments—Match the arguments of the request against the selected regular expression or regular expression class.
   - Request Body Length—Match packets where the body of the request is greater than the specified number of bytes.
   - Request Body—Match the body of the request against the selected regular expression or regular expression class.
   - Request Header Field Count—Match packets where the number of header fields in the request is greater than the specified count. You can match the field header type to a regular expression or to a predefined type. The predefined types are: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

- Request Header Field Length—Match packets where the length of the header field in the request is greater than the specified bytes. You can match the field header type to a regular expression or to a predefined type. The predefined types are listed above for Request Header Field Count.
- Request Header Field—Match the content of the selected header field in the request against the selected regular expression or regular expression class. You can specify a predefined header type or use a regular expression to select the headers.
- Request Header Count—Match packets where the number of headers in the request is greater than the specified number.
- Request Header Length—Match packets where the length of the header in the request is greater than the specified bytes.
- Request Header Non-ASCII—Match packets where the header in the request contains non-ASCII characters.
- Request Method—Match packets where the request method matches the predefined type or the selected regular expression or regular expression class. The predefined types are: bcopy, bdelete, bmove, bpropfind, bproppatch, connect, copy, delete, edit, get, getattribute, getattributenames, getproperties, head, index, lock, mkcol, mkdir, move, notify, options, poll, post, propfind, proppatch, put, revadd, revlabel, revlog, revnum, save, search, setattribute, startrev, stoprev, subscribe, trace, unedit, unlock, unsubscribe.
- Request URI Length—Match packets where the length of the URI of the request is greater than the specified bytes.
- Request URI—Match the content of the URI of the request against the selected regular expression or regular expression class.
- Request Body—Match the body of the request against the selected regular expression or regular expression class, or to ActiveX or Java Applet content.
- Response Body Length—Match packets where the length of the body of the response is greater than the specified bytes.
- Response Header Field Count—Match packets where the number of header fields in the response is greater than the specified count. You can match the field header type to a regular expression or to a predefined type. The predefined types are: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.
- Response Header Field Length—Match packets where the length of the header field in the response is greater than the specified bytes. You can match the field header type to a regular expression or to a predefined type. The predefined types are listed above for Response Header Field Count.
- Response Header Field—Match the content of the selected header field in the response against the selected regular expression or regular expression class. You can specify a predefined header type or use a regular expression to select the headers.
- Response Header Count—Match packets where the number of headers in the response is greater than the specified number.
- Response Header Length—Match packets where the length of the header in the response is greater than the specified bytes.

- Response Header Non-ASCII—Match packets where the header in the response contains non-ASCII characters.
- Response Status Line—Match the content of the response status line against the selected regular expression or regular expression class.

d) Choose whether to drop the connection, reset it, or log it. For drop connection and reset, you can enable or disable logging.
e) Click **OK** to add the inspection. Repeat the process as needed.

Step 7 Click **OK** in the HTTP Inspect Map dialog box.

You can now use the inspection map in a HTTP inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

ICMP Inspection

The ICMP inspection engine allows ICMP traffic to have a "session" so it can be inspected like TCP and UDP traffic. Without the ICMP inspection engine, we recommend that you do not allow ICMP through the ASA in an ACL. 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.

However, ICMP traffic directed to an ASA interface is never inspected, even if you enable ICMP inspection. Thus, a ping (echo request) to an interface can fail under specific circumstances, such as when the echo request comes from a source that the ASA can reach through a backup default route.

Note NAT uses ICMP inspection when translating packets even if you disable ICMP inspection.

For information on enabling ICMP inspection, see Configure Application Layer Protocol Inspection, on page 306.

ICMP Error Inspection

When ICMP Error inspection is enabled, the ASA creates translation sessions for intermediate hops that send ICMP error messages, based on the NAT configuration. The ASA overwrites the packet with the translated IP addresses.

When disabled, the ASA does not create translation sessions for intermediate nodes that generate ICMP error messages. ICMP error messages generated by the intermediate nodes between the inside host and the ASA reach the outside host without consuming any additional NAT resource. This is undesirable when an outside host uses the traceroute command to trace the hops to the destination on the inside of the ASA. When the ASA does not translate the intermediate hops, all the intermediate hops appear with the mapped destination IP address.

Note You should always enable ICMP Error inspection if there is a possibility that NAT will be used on ICMP packets. Because NAT automatically uses ICMP inspection for ICMP packets, even if you have ICMP inspection disabled, the use of the mapped destination address as the source address can make it look like a scanner is examining your network. For example, without ICMP Error inspection also enabled, if the echo request packet has its destination translated, when it is embedded in a ICMP time exceeded response, the outer header of the time exceeded request uses the translated destination as the source address. If you enable ICMP Error inspection, the time exceeded source address will be set to the correct value.

For information on enabling ICMP Error inspection, see Configure Application Layer Protocol Inspection, on page 306.

ILS Inspection

The Internet Locator Service (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. You cannot use PAT with ILS inspection because only IP addresses are stored by an LDAP database.

For search responses, when the LDAP server is located outside, consider using NAT to allow internal peers to communicate locally while registered to external LDAP servers. If you do not need to use NAT, we recommend that you turn off the inspection engine to provide better performance.

Additional configuration may be necessary when the ILS server is located inside the ASA border. This would require a hole for outside clients to access the LDAP server on the specified port, typically TCP 389.

Note Because ILS traffic (H225 call signaling) 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 TCP **timeout** command. In ASDM, this is on the **Configuration > Firewall > Advanced > Global Timeouts** pane.

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.

For information on enabling ILS inspection, see Configure Application Layer Protocol Inspection, on page 306.

Instant Messaging Inspection

The Instant Messaging (IM) inspect engine lets you control the network usage of IM and stop leakage of confidential data, propagation of worms, and other threats to the corporate network.

IM inspection is not enabled in the default inspection policy, so you must enable it if you need this inspection. However, the default inspect class does include the default IM ports, so you can simply edit the default global inspection policy to add IM inspection. You can alternatively create a new service policy as desired, for example, an interface-specific policy.

If you decide to implement IM inspection, you can also configure an IM inspection policy map to specify actions when a message violates a parameter. The following procedure explains IM inspection policy maps.

You can optionally create an IM inspection class map to define the traffic class for IM inspection. The other option is to define the traffic classes directly in the IM inspection policy map. The difference between creating a class map and defining the traffic match directly in the inspection map is that you can create more complex match criteria and you can reuse class maps. This procedure explains inspection maps, but class maps are essentially the same, except that you do not specify the actions for matching traffic. You can configure IM class maps by selecting **Configuration > Firewall > Objects > Class Maps > Instant Messaging (IM)**.



Tip You can configure inspection maps while creating service policies, in addition to the procedure explained below. The contents of the map are the same regardless of how you create it.

Before you begin

Some traffic matching options use regular expressions for matching purposes. If you intend to use one of those techniques, first create the regular expression or regular expression class map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **Instant Messaging (IM)**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map and click **Edit**.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 Define the specific inspections you want to implement based on traffic characteristics.

You can define traffic matching criteria based on IM class maps, by configuring matches directly in the inspection map, or both.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose **Single Match** to define the criterion directly, or **Multiple Match**, in which case you select the IM class map that defines the criteria. Click **Manage** to create new class maps.

c) If you are defining the criterion here, choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). For example, if No Match is selected

on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map. Then, configure the criterion.

- Protocol—Match traffic of a specific IM protocol, such as Yahoo Messenger or MSN Messenger.
- Service—Match a specific IM service, such as chat, file transfer, web cam, voice chat, conference, or games.
- Version—Match the version of the IM message against the selected regular expression or regular expression class.
- Client Login Name—Match the source client login name of the IM message against the selected regular expression or regular expression class.
- Client Peer Login Name—Match the destination peer login name of the IM message against the selected regular expression or regular expression class.
- Source IP Address—Match the source IP address and mask.
- Destination IP Address—Match the destination IP address and mask.
- Filename—Match the filename of the IM message against the selected regular expression or regular expression class.

d) Choose whether to drop the connection, reset it, or log it. For drop connection and reset, you can enable or disable logging.
e) Click **OK** to add the inspection. Repeat the process as needed.

Step 5 Click **OK** in the IM Inspect Map dialog box.

You can now use the inspection map in a IM inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

IP Options Inspection

You can configure IP Options inspection to control which IP packets are allowed based on the contents of the IP Options field in the packet header. You can drop packets that have unwanted options, clear the options (and allow the packet), or allow the packet without change.

IP options provide control functions that are required in some situations but unnecessary for most common communications. In particular, IP options include provisions for time stamps, security, and special routing. Use of IP Options is optional, and the field can contain zero, one, or more options.

For a list of IP options, with references to the relevant RFCs, see the IANA page, http://www.iana.org/assignments/ip-parameters/ip-parameters.xhtml.

IP options inspection is enabled by default, but for RSVP traffic only. You need to configure it only if you want to allow additional options than the default map allows, or if you want to apply it to other types of traffic by using a non-default inspection traffic class map.

Note IP options inspection does not work on fragmented packets. For example, options are not cleared from fragments.

The following sections describe IP Options inspection.

Defaults for IP Options Inspection

IP Options inspection is enabled by default for RSVP traffic only, using the _default_ip_options_map inspection policy map.

- The Router Alert option is allowed.

  This option notifies transit routers to inspect the contents of the packet even when the packet is not destined for that router. This inspection is valuable when implementing RSVP and similar protocols that require relatively complex processing from the routers along the packet's delivery path. Dropping RSVP packets containing the Router Alert option can cause problems in VoIP implementations.

- Packets that contain any other options are dropped.

  Each time a packet is dropped due to inspection, syslog 106012 is issued. The message shows which option caused the drop. Use the **show service-policy inspect ip-options** command to view statistics for each option.

Configure an IP Options Inspection Policy Map

If you want to perform non-default IP options inspection, create an IP options inspection policy map to specify how you want to handle each option type.

Tip You can configure inspection maps while creating service policies, in addition to the procedure explained below. The contents of the map are the same regardless of how you create it.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **IP Options**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map and click **Edit**.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 Choose which options you want to allow by moving them from the Drop list to the Allow list.

Consider the following tips:

- The “default” option sets the default behavior for options not included in the map. If you move it to the Allowed list, even options shown in the Drop list will be allowed.
- For any option you allow, you can check the Clear box to remove the option from the packet header before transmitting the packet.
- Some options are listed by option type number. The number is the whole option type octet (copy, class, and option number), not just the option number portion of the octet. These option types might not represent real options. Non-standard options must be in the expected type-length-value format defined in the Internet Protocol RFC 791, http://tools.ietf.org/html/rfc791.
- If a packet includes more than one type of option, it is dropped so long as the action for one of those types is to drop the packet.

For a list of IP options, with references to the relevant RFCs, see the IANA page, http://www.iana.org/assignments/ip-parameters/ip-parameters.xhtml.

Step 5 Click **OK**.

You can now use the inspection map in an IP options inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

IPsec Pass Through Inspection

IPsec Pass Through inspection is not enabled in the default inspection policy, so you must enable it if you need this inspection. However, the default inspect class does include the default IPsec ports, so you can simply edit the default global inspection policy to add IPsec inspection. You can alternatively create a new service policy as desired, for example, an interface-specific policy.

The following sections describe the IPsec Pass Through inspection engine.

IPsec Pass Through Inspection Overview

Internet Protocol Security (IPsec) is a protocol suite for securing IP communications by authenticating and encrypting each IP packet of a data stream. IPsec also includes protocols for establishing mutual authentication between agents at the beginning of the session and negotiation of cryptographic keys to be used during the session. IPsec can be used to protect data flows between a pair of hosts (for example, computer users or servers), between a pair of security gateways (such as routers or firewalls), or between a security gateway and a host.

IPsec Pass Through application inspection provides convenient traversal of ESP (IP protocol 50) and AH (IP protocol 51) traffic associated with an IKE UDP port 500 connection. It avoids lengthy ACL configuration to permit ESP and AH traffic and also provides security using timeout and max connections.

Configure a policy map for IPsec Pass Through to specify the restrictions for ESP or AH traffic. You can set the per client max connections and the idle timeout.

NAT and non-NAT traffic is permitted. However, PAT is not supported.

Configure an IPsec Pass Through Inspection Policy Map

An IPsec Pass Through map lets you change the default configuration values used for IPsec Pass Through application inspection. You can use an IPsec Pass Through map to permit certain flows without using an ACL.

The configuration includes a default map, _default_ipsec_passthru_map, that sets no maximum limit on ESP connections per client, and sets the ESP idle timeout at 10 minutes. You need to configure an inspection policy map only if you want different values, or if you want to set AH values.

Tip You can configure inspection maps while creating service policies, in addition to the procedure explained below. The contents of the map are the same regardless of how you create it.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **IPsec Pass Through**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map to view its contents. You can change the security level directly, or click **Customize** to edit the map. The remainder of the procedure assumes you are customizing or adding a map.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 In the **Security Level** view of the IPsec Pass Through Inspect Map dialog box, select the level that best matches your desired configuration.

If one of the preset levels matches your requirements, you are now done. Just click **OK**, skip the rest of this procedure, and use the map in a service policy rule for IPsec Pass Through inspection.

If you need to customize the settings further, click **Details**, and continue with the procedure.

Step 5 Choose whether to allow ESP and AH tunnels.

For each protocol, you can also set the maximum number of connections allowed per client, and the idle timeout.

Step 6 Click **OK**.

You can now use the inspection map in an IPsec Pass Through inspection service policy.

IPv6 Inspection

IPv6 inspection lets you selectively log or drop IPv6 traffic based on the extension header. In addition, IPv6 inspection can check conformance to RFC 2460 for type and order of extension headers in IPv6 packets.

IPv6 inspection is not enabled in the default inspection policy, so you must enable it if you need this inspection. You can simply edit the default global inspection policy to add IPv6 inspection. You can alternatively create a new service policy as desired, for example, an interface-specific policy.

Defaults for IPv6 Inspection

If you enable IPv6 inspection and do not specify an inspection policy map, then the default IPv6 inspection policy map is used, and the following actions are taken:

- Allows only known IPv6 extension headers. Non-conforming packets are dropped and logged.
- Enforces the order of IPv6 extension headers as defined in the RFC 2460 specification. Non-conforming packets are dropped and logged.
- Drops any packet with a routing type header.

Configure an IPv6 Inspection Policy Map

To identify extension headers to drop or log, or to disable packet verification, create an IPv6 inspection policy map to be used by the service policy.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **IPv6**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map and click **Edit**.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 Click the **Enforcement** tab and choose whether to permit only known IPv6 extension headers or to enforce the order of IPv6 extension headers as defined in RFC 2460. Non-conforming packets are dropped and logged.

Step 5 (Optional) Click the **Header Matches** tab to identify traffic to drop or log based on the headers in IPv6 messages.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose the IPv6 extension header to match:

- Authentication (AH) header.
- Destination Options header.
- Encapsulating Security Payload (ESP) header.
- Fragment header.
- Hop-by-Hop Options header.
- Routing header—Specify either a single header type number or a range of numbers.
- Header count—Specify the maximum number of extension headers you will allow without dropping or logging the packet.
- Routing header address count—Specify the maximum number of addresses in the type 0 routing header you will allow without dropping or logging the packet.

c) Choose whether to drop or log the packet. If you drop the packet, you can also enable logging.
d) Click **OK** to add the inspection. Repeat the process as needed.

Step 6 Click **OK** in the IPv6 Inspect Map dialog box.

You can now use the inspection map in an IPv6 inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

NetBIOS Inspection

NetBIOS application inspection performs NAT for the embedded IP address in the NetBIOS name service (NBNS) packets and NetBIOS datagram services packets. It also enforces protocol conformance, checking the various count and length fields for consistency.

NetBIOS inspection is enabled by default. You can optionally create a policy map to drop or log NetBIOS protocol violations. The following procedure explains how to configure a NetBIOS inspection policy map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **NetBIOS**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map and click **Edit**.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 Select **Check for Protocol Violations**. There is no reason to create a map if you do not select this option.

Step 5 Select the action to take, either to drop the packet or log it. If you drop the packet, you can also enable logging.

Step 6 Click **OK**.

You can now use the inspection map in a NetBIOS inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

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 carry 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.

Specifically, the ASA inspects the PPTP version announcements and the outgoing call request/response sequence. Only PPTP Version 1, as defined in RFC 2637, is inspected. Further inspection on the TCP control channel is disabled if the version announced by either side is not Version 1. In addition, the outgoing-call request and reply sequence are tracked. Connections and xlates are dynamically 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).

For information on enabling PPTP inspection, see Configure Application Layer Protocol Inspection, on page 306.

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.

For information on enabling RSH inspection, see Configure Application Layer Protocol Inspection, on page 306.

SMTP and Extended SMTP Inspection

ESMTP inspection detects attacks, including spam, phising, malformed message attacks, and buffer overflow/underflow attacks. It also provides support for application security and protocol conformance, which enforces the sanity of the ESMTP messages as well as block senders/receivers, and block mail relay.

ESMTP inspection is enabled by default. You need to configure it only if you want different processing than that provided by the default inspection map.

The following sections describe the ESMTP inspection engine.

SMTP and ESMTP Inspection Overview

Extended SMTP (ESMTP) application inspection provides improved protection against SMTP-based attacks by restricting the types of SMTP commands that can pass through the ASA and by adding monitoring capabilities. ESMTP is an enhancement to the SMTP protocol and is similar is most respects to SMTP.

ESMTP application inspection controls and reduces the commands that the user can use as well as the messages that the server returns. ESMTP inspection performs three primary tasks:

- Restricts SMTP requests to seven basic SMTP commands and eight extended commands. Supported commands are the following:
  - Extended SMTP—AUTH, EHLO, ETRN, HELP, SAML, SEND, SOML, STARTTLS, and VRFY.
  - SMTP (RFC 821)—DATA, HELO, MAIL, NOOP, QUIT, RCPT, RSET.
- Monitors the SMTP command-response sequence.
- Generates an audit trail—Audit record 108002 is generated when an invalid character embedded in the mail address is replaced. For more information, see RFC 821.

ESMTP 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 or unsupported commands, the inspection engine changes all the characters in the packet to X, which are rejected by the internal server. This results in a message such as “500 Command unknown: 'XXX'.” Incomplete commands are discarded

  Unsupported ESMTP commands are ATRN, ONEX, VERB, CHUNKING, and private extensions..
- TCP stream editing.
- Command pipelining.



Note With ESMTP 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; they must be terminated with carriage return and line feed; and you must wait for a response before issuing the next reply.

Defaults for ESMTP Inspection

ESMTP inspection is enabled by default, using the _default_esmtp_map inspection policy map.

- The server banner is masked. 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.
- Encrypted connections are allowed but not inspected.
- Special characters in sender and receiver address are not noticed, no action is taken.
- Connections with command line length greater than 512 are dropped and logged.
- Connections with more than 100 recipients are dropped and logged.
- Messages with body length greater than 998 bytes are logged.
- Connections with header line length greater than 998 are dropped and logged.
- Messages with MIME filenames greater than 255 characters are dropped and logged.
- EHLO reply parameters matching “others” are masked.

Configure an ESMTP Inspection Policy Map

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.

Before you begin

Some traffic matching options use regular expressions for matching purposes. If you intend to use one of those techniques, first create the regular expression or regular expression class map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **ESMTP**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map to view its contents. You can change the security level directly, or click **Customize** to edit the map. The remainder of the procedure assumes you are customizing or adding a map.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 In the **Security Level** view of the ESMTP Inspect Map dialog box, select the level that best matches your desired configuration.

If one of the preset levels matches your requirements, you are now done. Just click **OK**, skip the rest of this procedure, and use the map in a service policy rule for ESMTP inspection.

If you need to customize the settings further, click **Details**, and continue with the procedure.

Tip
The **MIME File Type Filtering** button is a shortcut to configure file type inspection, which is explained later in this procedure.

Step 5 Click the **Parameters** tab and configure the desired options.

- **Mask Server Banner**—Whether to mask the banner from the ESMTP server.
- **Encrypted Packet Inspection**—Whether to allow ESMTP over TLS (encrypted connections) without inspection. You can optionally log encrypted connections. The default is to allow TLS sessions without inspection. If you deselect the option, the system strips the STARTTLS indication from any encrypted session connection attempt and forces a plain-text connection.

Step 6 Click the **Filtering** tab and configure the desired options.

- **Configure mail relay**—Identifies a domain name for mail relay. You can either drop the connection and optionally log it, or log it.
- **Check for special characters**—Identifies the action to take for messages that include the special characters pipe (|), back quote, and NUL in the sender or receiver email addresses. You can either drop the connection and optionally log it, or log it.

Step 7 Click the **Inspections** tab and define the specific inspections you want to implement based on traffic characteristics.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map. Then, configure the criterion:

- **Body Length**—Matches messages where the length of an ESMTP body message is greater than the specified number of bytes.
- **Body Line Length**—Matches messages where the length of a line in an ESMTP body message is greater than the specified number of bytes.
- **Commands**—Matches the command verb in the message. You can specify one or more of the following commands: auth, data, ehlo, etrn, helo, help, mail, noop, quit, rcpt, rset, saml, soml, vrfy.
- **Command Recipient Count**—Matches messages where the number of recipients is greater than the specified count.
- **Command Line Length**—Matches messages where the length of a line in the command verb is greater than the specified number of bytes.
- **EHLO Reply Parameters**—Matches ESMTP EHLO reply parameters. You can specify one or more of the following parameters: 8bitmime, auth, binaryname, checkpoint, dsn, etrn, others, pipelining, size, vrfy.
- **Header Length**—Matches messages where the length of an ESMTP header is greater than the specified number of bytes.

- **Header Line Length**—Matches messages where the length of a line in an ESMTP header is greater than the specified number of bytes.
- **Header To: Fields Count**—Matches messages where the number of To fields in the header is greater than the specified number.
- **Invalid Recipients Count**—Matches messages where the number of invalid recipients is greater than the specified count.
- **MIME File Type**—Matches the MIME or media file type against the specified regular expression or regular expression class.
- **MIME Filename Length**—Matches messages where a file name is longer than the specified number of bytes.
- **MIME Encoding**—Matches the MIME encoding type. You can specify one or more of the following types: 7bit, 8bit, base64, binary, others, quoted-printable.
- **Sender Address**—Matches the sender email address against the specified regular expression or regular expression class.
- **Sender Address Length**—Matches messages where the sender address is greater than the specified number of bytes.

c) Choose whether to drop the connection, reset it, or log it. For drop connection and reset, you can enable or disable logging. For command and EHLO reply parameter matching, you can also mask the command. For command matching, you can also apply a rate limit in packets per second.
d) Click **OK** to add the inspection. Repeat the process as needed.

Step 8 Click **OK** in the ESMTP Inspect Map dialog box.

You can now use the inspection map in a ESMTP inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

SNMP Inspection

SNMP application inspection is applied to both to-the-device and through-the-device traffic. This inspection is necessary if you configure SNMP v3 where users are limited to specific SNMP hosts. Without the inspection, a defined v3 user can poll the device from any allowed host. SNMP inspection is enabled by default for the default ports, so you need to configure it only if you use non-default ports. The default ports are UDP/161, 162 (for all device types) and UDP/4161 for devices that also run FXOS, as FXOS listens on UDP/161.

By default, the SNMP inspection limits the polling to the configured version.

Note If you configure SNMP on a device that also runs FXOS, SNMP inspection is mandatory, and is re-enabled if you disable it. SNMP inspection is enabled on a traffic class map that includes port UDP/4161.

Optionally, you can further 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 system can deny SNMP versions 1, 2, 2c, or 3. You control the versions permitted by creating an SNMP map, as explained below. If you do not need to control the versions, simply enable SNMP inspection without a map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **SNMP**.

Step 2 Click **Add**, or select a map and click **Edit**. When adding a map, enter a map name.

Step 3 Select the SNMP versions to disallow.

Step 4 Click **OK**.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

SQL*Net Inspection

SQL*Net inspection is enabled by default. The inspection engine supports SQL*Net versions 1 and 2, but only the Transparent Network Substrate (TNS) format. Inspection does not support the Tabular Data Stream (TDS) format. SQL*Net messages are scanned for embedded addresses and ports, and NAT rewrite is applied when necessary.

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). If your application uses a different port, apply the SQL*Net inspection to a traffic class that includes that port.

Disable SQL*Net inspection when:

- SQL data transfer occurs on the same port as the SQL control TCP port 1521. The security appliance acts as a proxy when SQL*Net inspection is enabled and reduces the client window size from 65000 to about 16000 causing data transfer issues.
- The inspection of high rates of SQL traffic causes unacceptable spikes in CPU usage.

After disabling SQL*Net inspection, use the **clear conn port 1521** command so that connections can be rebuilt without inspection.

For information on enabling SQL*Net inspection, see Configure Application Layer Protocol Inspection, on page 306.

Sun RPC Inspection

This section describes Sun RPC application inspection.

Sun RPC Inspection Overview

Sun RPC protocol inspection is enabled by default. You simply need to manage the Sun RPC server table to identify which services are allowed to traverse the firewall. However, pinholing for NFS is done for any server even without the server table configuration.

Sun RPC is used by NFS and NIS. Sun RPC services can run on any port. When a client attempts to access a Sun RPC service on a server, it must learn the port that service is running on. It does this by querying the port mapper process, usually rpcbind, on the well-known port of 111.

The client sends the Sun RPC program number of the service and the port mapper process responds with the port number of the service. The client sends its Sun RPC queries to the server, specifying the port identified by the port mapper process. When the server replies, the ASA intercepts this packet and opens both embryonic TCP and UDP connections on that port.

NAT or PAT of Sun RPC payload information is not supported.

Manage Sun RPC Services

Use the Sun RPC services table to control Sun RPC traffic based on established Sun RPC sessions.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Advanced** > **SUNRPC Server**.

Step 2 Do one of the following:

- Click **Add** to add a new server.
- Select a server and click **Edit**.

Step 3 Configure the service properties:

- **Interface Name**—The interface through which traffic to the server flows.
- **IP Address/Mask**—The address of the Sun RPC server.
- **Service ID**—The service type on the server. To determine the service type (for example, 100003), use the **sunrpcinfo** command at the UNIX or Linux command line on the Sun RPC server machine.
- **Protocol**—Whether the service uses TCP or UDP.
- **Port/Port Range**—The port or range of ports used by the service.
- **Timeout**—The idle timeout for the pinhole opened for the connection by Sun RPC inspection.

Step 4 Click **OK**.

Step 5 (Optional.) Monitor the pinholes created for these services.

To display the pinholes open for Sun RPC services, enter the **show sunrpc-server active** command. Select **Tools** > **Command Line Interface** to enter the command. For example:

```
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.

If necessary, you can clear these services using the **clear sunrpc-server active**

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 inspection engine inspects TFTP read request (RRQ), write request (WRQ), and error notification (ERROR), and dynamically creates connections and translations, if necessary, to permit file transfer between a TFTP client and server.

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.

For information on enabling TFTP inspection, see Configure Application Layer Protocol Inspection, on page 306.

XDMCP Inspection

XDMCP is a protocol that uses UDP port 177 to negotiate X sessions, which use TCP when established.

For successful negotiation and start of an XWindows session, the ASA must allow the TCP back connection from the Xhosted computer. To permit the back connection, you can use access rules to allow the TCP ports. Alternatively, you can use the **established** command on the ASA. Once XDMCP negotiates the port to send the display, the **established** command is consulted to verify if this back connection should be permitted.

During the XWindows session, the manager talks to the display Xserver on the well-known port 6000 | n. Each display has a separate connection to the Xserver, as a result of the following terminal setting.

```
setenv DISPLAY Xserver:n
```

where *n* is the display number.

When XDMCP is used, the display is negotiated using IP addresses, which the ASA can NAT if needed. XDCMP inspection does not support PAT.

For information on enabling XDMCP inspection, see Configure Application Layer Protocol Inspection, on page 306.

VXLAN Inspection

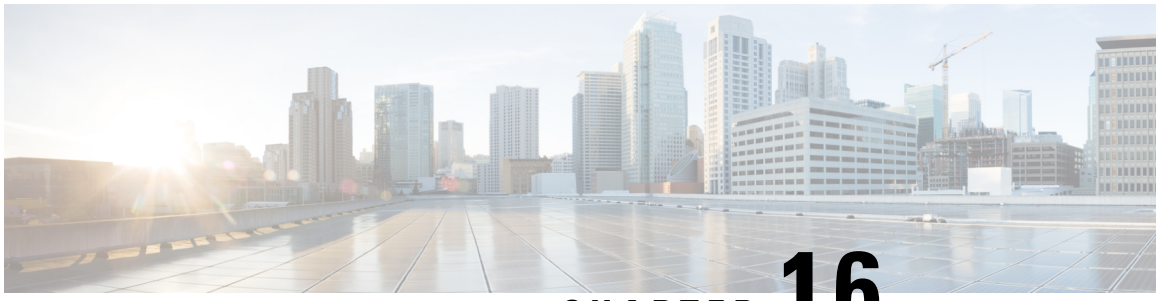
Virtual Extensible Local Area Network (VXLAN) inspection works on VXLAN encapsulated traffic that passes through the ASA. It ensures that the VXLAN header format conforms to standards, dropping any malformed packets. VXLAN inspection is not done on traffic for which the ASA acts as a VXLAN Tunnel End Point (VTEP) or a VXLAN gateway, as those checks are done as a normal part of decapsulating VXLAN packets.

VXLAN packets are UDP, normally on port 4789. This port is part of the default-inspection-traffic class, so you can simply add VXLAN inspection to the inspection_default service policy rule. Alternatively, you can create a class for it using port or ACL matching.

History for Basic Internet Protocol Inspection

Feature Name	Releases	Feature Information
DCERPC inspection support for ISystemMapper UUID message RemoteGetClassObject opnum3.	9.4(1)	The ASA started supporting non-EPM DCERPC messages in release 8.3, supporting the ISystemMapper UUID message RemoteCreateInstance opnum4. This change extends support to the RemoteGetClassObject opnum3 message. We did not modify any ASDM screens.
VXLAN packet inspection	9.4(1)	The ASA can inspect the VXLAN header to enforce compliance with the standard format. We modified the following screen: **Configuration > Firewall > Service Policy Rules > Add Service Policy Rule > Rule Actions > Protocol Inspection**.
ESMTP inspection change in default behavior for TLS sessions.	9.4(1)	The default for ESMTP inspection was changed to allow TLS sessions, which are not inspected. However, this default applies to new or reimaged systems. If you upgrade a system that includes **no allow-tls**, the command is not changed. The change in default behavior was also made in these older versions: 8.4(7.25), 8.5(1.23), 8.6(1.16), 8.7(1.15), 9.0(4.28), 9.1(6.1), 9.2(3.2) 9.3(1.2), 9.3(2.2).
IP Options inspection improvements.	9.5(1)	IP Options inspection now supports all possible IP options. You can tune the inspection to allow, clear, or drop any standard or experimental options, including those not yet defined. You can also set a default behavior for options not explicitly defined in an IP options inspection map. We changed the IP Options Inspect Map dialog box to include additional options. You now select which options to allow and optionally clear.

Feature Name	Releases	Feature Information
DCERPC inspection improvements and UUID filtering	9.5(2)	DCERPC inspection now supports NAT for OxidResolver ServerAlive2 opnum5 messages. You can also now filter on DCERPC message universally unique identifiers (UUIDs) to reset or log particular message types. There is a new DCERPC inspection class map for UUID filtering. We added the following screen: **Configuration** > **Firewall** > **Objects** > **Class Maps** > **DCERPC**. We modified the following screen: **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **DCERPC**.
DNS over TCP inspection.	9.6(2)	You can now inspect DNS over TCP traffic (TCP/53). We modified the following page: **Configuration** > **Firewall** > **Objects** > **Inspection Maps** > **DNS** Add/Edit dialog box.
Cisco Umbrella support.	9.10(1)	You can configure the device to redirect DNS requests to Cisco Umbrella, so that your Enterprise Security policy defined in Cisco Umbrella can be applied to user connections. You can allow or block connections based on FQDN, or for suspicious FQDNs, you can redirect the user to the Cisco Umbrella intelligent proxy, which can perform URL filtering. The Umbrella configuration is part of the DNS inspection policy. We added or modified the following screens: **Configuration** > **Firewall** > **Objects** > **Umbrella**, **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **DNS**.
Cisco Umbrella Enhancements.	9.12(1)	You can now identify local domain names that should bypass Cisco Umbrella. DNS requests for these domains go directly to the DNS servers without Umbrella processing. You can also identify which Umbrella servers to use for resolving DNS requests. Finally, you can define the Umbrella inspection policy to fail open, so that DNS requests are not blocked if the Umbrella server is unavailable. We modified the following screens: **Configuration** > **Firewall** > **Objects** > **Umbrella**, **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **DNS**.
XDMCP inspection disabled by default in new installations.	9.15(1)	Previously, XDMCP inspection was enabled by default for all traffic. Now, on new installations, which includes new systems and reimaged systems, XDMCP is off by default. If you need this inspection, please enable it. Note that on upgrades, your current settings for XDMCP inspection are retained, even if you simply had it enabled by way of the default inspection settings.



CHAPTER 16

Inspection for Voice and Video Protocols

The following topics explain application inspection for voice and video protocols. For basic information on why you need to use inspection for certain protocols, and the overall methods for applying inspection, see Getting Started with Application Layer Protocol Inspection, on page 299.

- CTIQBE Inspection, on page 351
- H.323 Inspection, on page 352
- MGCP Inspection, on page 356
- RTSP Inspection, on page 359
- SIP Inspection, on page 361
- Skinny (SCCP) Inspection, on page 366
- STUN Inspection, on page 369
- History for Voice and Video Protocol Inspection, on page 370

CTIQBE Inspection

CTIQBE protocol inspection supports NAT, PAT, and bidirectional NAT. This enables Cisco IP SoftPhone and other Cisco TAPI/JTAPI applications to work successfully with Cisco CallManager for call setup across the ASA.

TAPI and JTAPI are used by many Cisco VoIP applications. CTIQBE is used by Cisco TSP to communicate with Cisco CallManager.

For information on enabling CTIQBE inspection, see Configure Application Layer Protocol Inspection, on page 306.

Limitations for CTIQBE Inspection

Stateful failover of CTIQBE calls is not supported.

The following summarizes special considerations when using CTIQBE application inspection in specific scenarios:

- If two Cisco IP SoftPhones are registered with different Cisco CallManagers, which are connected to different interfaces of the ASA, calls between these two phones fail.
- When Cisco CallManager is located on the higher security interface compared to Cisco IP SoftPhones, if NAT or outside NAT is required for the Cisco CallManager IP address, the mapping must be static as

Cisco IP SoftPhone requires the Cisco CallManager IP address to be specified explicitly in its Cisco TSP configuration on the PC.

- When using PAT or Outside PAT, if the Cisco CallManager IP address is to be translated, its TCP port 2748 must be statically mapped to the same port of the PAT (interface) address for Cisco IP SoftPhone registrations to succeed. The CTIQBE listening port (TCP 2748) is fixed and is not user-configurable on Cisco CallManager, Cisco IP SoftPhone, or Cisco TSP.

H.323 Inspection

H.323 inspection supports RAS, H.225, and H.245, and its functionality translates all embedded IP addresses and ports. It performs state tracking and filtering and can do a cascade of inspect function activation. H.323 inspection supports phone number filtering, dynamic T.120 control, H.245 tunneling control, HSI groups, protocol state tracking, H.323 call duration enforcement, T.38 Fax, and audio/video control.

H.323 inspection is enabled by default. You need to configure it only if you want non-default processing.

The following sections describe the H.323 application inspection.

H.323 Inspection Overview

H.323 inspection provides support for H.323 compliant applications such as Cisco CallManager. H.323 is a suite of protocols defined by the International Telecommunication Union for multimedia conferences over LANs. The ASA supports H.323 through Version 6, including H.323 v3 feature Multiple Calls on One Call Signaling Channel.

With H.323 inspection enabled, the ASA supports multiple calls on the same call signaling channel, a feature introduced with H.323 Version 3. This feature reduces call setup time and reduces the use of ports on the ASA.

The two major functions of H.323 inspection are as follows:

- NAT the necessary embedded IPv4 addresses in the H.225 and H.245 messages. Because H.323 messages are encoded in PER encoding format, the ASA uses an ASN.1 decoder to decode the H.323 messages.
- Dynamically allocate the negotiated H.245 and RTP/RTCP connections. The H.225 connection can also be dynamically allocated when using RAS.

How H.323 Works

The H.323 collection of protocols collectively may use up to two TCP connection and four to eight UDP connections. FastConnect uses only one TCP connection, and RAS uses a single UDP connection for registration, admissions, and status.

An H.323 client can initially establish a TCP connection to an H.323 server using TCP port 1720 to request Q.931 call setup. As part of the call setup process, the H.323 terminal supplies a port number to the client to use for an H.245 TCP connection. In environments where H.323 gatekeeper is in use, the initial packet is transmitted using UDP.

H.323 inspection monitors the Q.931 TCP connection to determine the H.245 port number. If the H.323 terminals are not using FastConnect, the ASA dynamically allocates the H.245 connection based on the inspection of the H.225 messages. The H.225 connection can also be dynamically allocated when using RAS.

Within each H.245 message, the H.323 endpoints exchange port numbers that are used for subsequent UDP data streams. H.323 inspection inspects the H.245 messages to identify these ports and dynamically creates connections for the media exchange. RTP uses the negotiated port number, while RTCP uses the next higher port number.

The H.323 control channel handles H.225 and H.245 and H.323 RAS. H.323 inspection uses the following ports.

- 1718—Gate Keeper Discovery UDP port
- 1719—RAS UDP port
- 1720—TCP Control Port

You must permit traffic for the well-known H.323 port 1719 for RAS signaling. Additionally, you must permit traffic for the well-known H.323 port 1720 for the H.225 call signaling; however, the H.245 signaling ports are negotiated between the endpoints in the H.225 signaling. When an H.323 gatekeeper is used, the ASA opens an H.225 connection based on inspection of the ACF and RCF messages.

After inspecting the H.225 messages, the ASA opens the H.245 channel and then inspects traffic sent over the H.245 channel as well. All H.245 messages passing through the ASA undergo H.245 application inspection, which translates embedded IP addresses and opens the media channels negotiated in H.245 messages.

Each UDP connection with a packet going through H.323 inspection is marked as an H.323 connection and times out with the H.323 timeout as configured in the Configuration > Firewall > Advanced > Global Timeouts pane.

Note You can enable call setup between H.323 endpoints when the Gatekeeper is inside the network. The ASA includes options to open pinholes for calls based on the RegistrationRequest/RegistrationConfirm (RRQ/RCF) messages. Because these RRQ/RCF messages are sent to and from the Gatekeeper, the calling endpoint's IP address is unknown and the ASA opens a pinhole through source IP address/port 0/0. By default, this option is disabled.

H.239 Support in H.245 Messages

The ASA sits between two H.323 endpoints. When the two H.323 endpoints set up a telepresentation session so that the endpoints can send and receive a data presentation, such as spreadsheet data, the ASA ensure successful H.239 negotiation between the endpoints.

H.239 is a standard that provides the ability for H.300 series endpoints to open an additional video channel in a single call. In a call, an endpoint (such as a video phone), sends a channel for video and a channel for data presentation. The H.239 negotiation occurs on the H.245 channel.

The ASA opens pinholes for the additional media channel and the media control channel. The endpoints use open logical channel message (OLC) to signal a new channel creation. The message extension is part of H.245 version 13.

The decoding and encoding of the telepresentation session is enabled by default. H.239 encoding and decoding is preformed by ASN.1 coder.

Limitations for H.323 Inspection

H.323 inspection is tested and supported for Cisco Unified Communications Manager (CUCM) 7.0. It is not supported for CUCM 8.0 and higher. H.323 inspection might work with other releases and products.

The following are some of the known issues and limitations when using H.323 application inspection:

- PAT is supported except for extended PAT or per-session PAT.
- 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.
- Not supported with NAT between same-security-level interfaces.
- Not supported with NAT64.
- NAT with H.323 inspection is not compatible with NAT when done directly on the endpoints. If you perform NAT on the endpoints, disable H.323 inspection.

Configure H.323 Inspection Policy Map

You can create an H.323 inspection policy map to customize H.323 inspection actions if the default inspection behavior is not sufficient for your network.

You can optionally create a H.323 inspection class map to define the traffic class for H.323 inspection. The other option is to define the traffic classes directly in the H.323 inspection policy map. The difference between creating a class map and defining the traffic match directly in the inspection map is that you can create more complex match criteria and you can reuse class maps. Although this procedure explains inspection maps, the matching criteria used in class maps are the same as those explained in the step relating to the **Inspection** tab. You can configure H.323 class maps by selecting **Configuration** > **Firewall** > **Objects** > **Class Maps** > **H.323**, or by creating them while configuring the inspection map.

Tip You can configure inspection maps while creating service policies, in addition to the procedure explained below. The contents of the map are the same regardless of how you create it.

Before you begin

Some traffic matching options use regular expressions for matching purposes. If you intend to use one of those techniques, first create the regular expression or regular expression class map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **H.323**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map to view its contents. You can change the security level directly, or click **Customize** to edit the map. The remainder of the procedure assumes you are customizing or adding a map.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 In the **Security Level** view of the H.323 Inspect Map dialog box, select the level that best matches your desired configuration. The default level is Low.

If one of the preset levels matches your requirements, you are now done. Just click **OK**, skip the rest of this procedure, and use the map in a service policy rule for H.323 inspection.

Tip
The **Phone Number Filtering** button is a shortcut to configure called or calling party inspection, which is explained later in this procedure.

Step 5 If you need to customize the settings further, click **Details**, and do the following:

a) Click the **State Checking** tab and choose whether to enable state transition checking of RAS and H.225 messages.

You can also check RCF messages and open pinholes for call signal addresses present in RRQ messages, which enables call setup between H.323 endpoints when the Gatekeeper is inside the network. Use this option to open pinholes for calls based on the RegistrationRequest/RegistrationConfirm (RRQ/RCF) messages. Because these RRQ/RCF messages are sent to and from the Gatekeeper, the calling endpoint's IP address is unknown and the ASA opens a pinhole through source IP address/port 0/0. By default, this option is disabled.

b) Click the **Call Attributes** tab and choose whether to enforce a call duration limit (maximum is 1193 hours) or to enforce the presence of calling and called party numbers during call setup.

You an also allow H.225 FACILITY messages to arrive before H.225 SETUP messages in accordance to H.460.18. If you encounter call setup issues, where connections are being closed before being completed when using H.323/H.225, select this option to allow early messages. Also, ensure that you enable inspection for both H.323 RAS and H.225 (they are both enabled by default).

c) Click the **Tunneling and Protocol Conformance** tab and choose whether check for H.245 tunneling; you can either drop the connection or log it.

You can also choose whether to check RTP packets that are flowing on the pinholes for protocol conformance. If you check for conformance, you can also choose whether to limit the payload to audio or video, based on the signaling exchange.

Step 6 If necessary, click the **HSI Group Parameters** tab and define the HSI groups.

a) Do any of the following:

- Click **Add** to add a new group.
- Select an existing group and click **Edit**.

b) Specify the group ID (from 0 to 2147483647) and the IP address of the HSI.

c) To add an endpoint to the HSI group, enter the IP address, select the interface through which the endpoint is connected to the ASA, and click **Add>>**. Remove any endpoints that are no longer needed. You can have up to 10 endpoints per group.

d) Click **OK** to add the group. Repeat the process as needed.

Step 7 Click the **Inspections** tab and define the specific inspections you want to implement based on traffic characteristics.

You can define traffic matching criteria based on H.323 class maps, by configuring matches directly in the inspection map, or both.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose **Single Match** to define the criterion directly, or **Multiple Match**, in which case you select the H.323 class map that defines the criteria.

c) If you are defining the criterion here, choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). Then, configure the criterion as follows:

- **Called Party**—Match the H.323 called party against the selected regular expression or regular expression class.
- **Calling Party**—Match the H.323 calling party against the selected regular expression or regular expression class.
- **Media Type**—Match the media type: audio, video, or data.

d) Choose the action to take for matching traffic. For calling or called party matching, you can drop the packet, drop the connection, or reset the connection. For media type matching, the action is always to drop the packet; you can enable logging for this action.

e) Click **OK** to add the inspection. Repeat the process as needed.

Step 8 Click **OK** in the H.323 Inspect Map dialog box.

You can now use the inspection map in an H.323 inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

MGCP Inspection

MGCP inspection is not enabled in the default inspection policy, so you must enable it if you need this inspection. However, the default inspect class does include the default MGCP ports, so you can simply edit the default global inspection policy to add MGCP inspection. You can alternatively create a new service policy as desired, for example, an interface-specific policy.

The following sections describe MGCP application inspection.

MGCP Inspection Overview

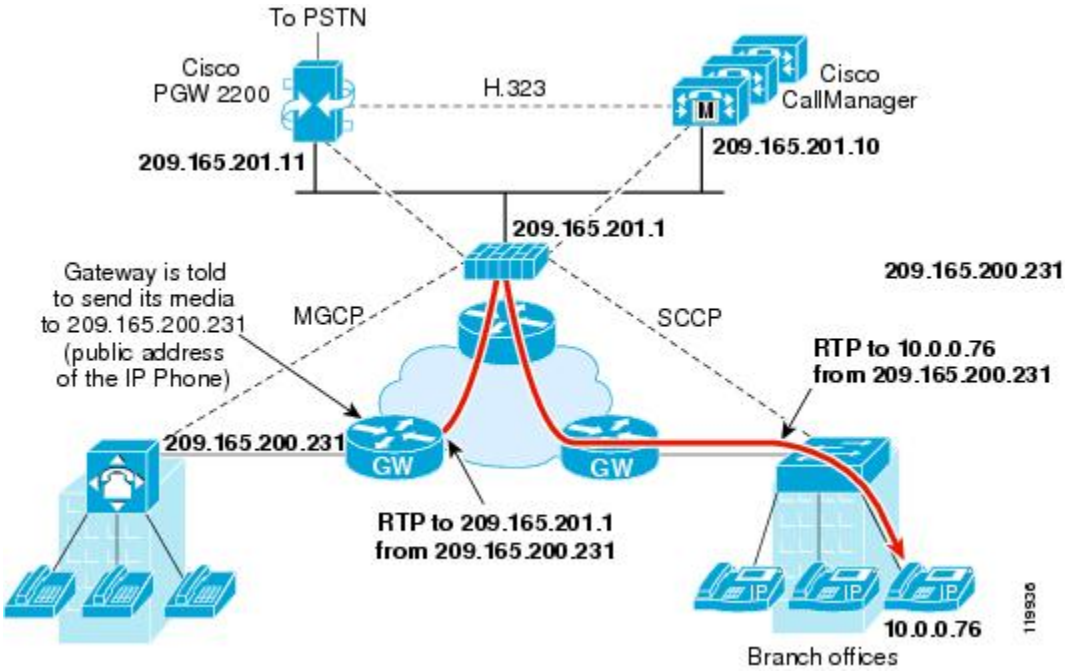
MGCP is used to control media gateways from external call control elements called media gateway controllers or call agents. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over 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.

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. The following figure illustrates how you can use NAT with MGCP.

Figure 47: 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.

- Gateways usually listen to UDP port 2427 to receive commands from the call agent.
- The port on which the call agent receives commands from the gateway. Call agents usually listen to UDP port 2727 to receive commands from the gateway.

Note MGCP inspection does not support the use of different IP addresses for MGCP signaling and RTP data. A common and recommended practice is to send RTP data from a resilient IP address, such as a loopback or virtual IP address; however, the ASA requires the RTP data to come from the same address as MGCP signaling.

Configure an MGCP Inspection Policy Map

If the network has multiple call agents and gateways for which the ASA has to open pinholes, create an MGCP map. You can then apply the MGCP map when you enable MGCP inspection.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **MGCP**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map and click **Edit**.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 (Optional) Click the **Command Queue** tab and specify the maximum number of commands allowed in the MGCP command queue. The default is 200, the allowed range is 1 to 2147483647.

Step 5 Click the **Gateways and Call Agents** tab and configure the groups of gateways and call agents for the map.

a) Click **Add** to create a new group, or select a group and click **Edit**.

b) Enter the **Group ID** of the call agent group. A call agent group associates one or more call agents with one or more MGCP media gateways. The valid range is from 0 to 2147483647.

c) Add the IP addresses of the media gateways that are controlled by the associated call agents to the group by entering them in **Gateway to Be Added** and clicking **Add>>**. Delete any gateways that are no longer used.

A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Normally, a gateway sends commands to the default MGCP port for call agents, UDP 2727.

d) Add the IP addresses of the call agents that control the MGCP media gateways by entering them in **Call Agent to Be Added** and clicking **Add>>**. Delete any agents that are no longer needed.

Normally, a call agent sends commands to the default MGCP port for gateways, UDP 2427.

e) Click **OK** in the MGCP Group dialog box. Repeat the process to add other groups as needed.

Step 6 Click **OK** in the MGCP Inspect Map dialog box.

You can now use the inspection map in an MGCP inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

RTSP Inspection

RTSP inspection is enabled by default. You need to configure it only if you want non-default processing. The following sections describe RTSP application inspection.

RTSP Inspection Overview

The RTSP inspection engine lets the ASA pass RTSP packets. RTSP is used by RealAudio, RealNetworks, Apple QuickTime 4, RealPlayer, and Cisco IP/TV connections.

Note For Cisco IP/TV, use RTSP TCP ports 554 and 8554.

RTSP applications use the well-known port 554 with TCP (rarely UDP) as a control channel. The ASA only supports TCP, in conformity with RFC 2326. This TCP control channel is used to negotiate the data channels that are used to transmit audio/video traffic, depending on the transport mode that is configured on the client.

The supported RDT transports are: rtp/avp, rtp/avp/udp, x-real-rdt, x-real-rdt/udp, and x-pn-tng/udp.

The ASA parses Setup response messages with a status code of 200. If the response message is traveling inbound, the server is outside relative to the ASA and dynamic channels need to be opened for connections coming inbound from the server. If the response message is outbound, then the ASA does not need to open dynamic channels.

RTSP inspection does not support PAT or dual-NAT. Also, the ASA cannot recognize HTTP cloaking where RTSP messages are hidden in the HTTP messages.

RealPlayer Configuration Requirements

When using RealPlayer, it is important to properly configure transport mode. For the ASA, add an **access-list** command from the server to the client or vice versa. For RealPlayer, change transport mode by clicking **Options**>**Preferences**>**Transport**>**RTSP Settings**.

If using TCP mode on the RealPlayer, select the **Use TCP to Connect to Server** and **Attempt to use TCP for all content** check boxes. On the ASA, there is no need to configure the inspection engine.

If using UDP mode on the RealPlayer, select the **Use TCP to Connect to Server** and **Attempt to use UDP for static content** check boxes, and for live content not available via multicast. On the ASA, add an **inspect rtsp** command.

Limitations for RSTP Inspection

The following restrictions apply to the RSTP inspection.

- The ASA does not support multicast RTSP or RTSP messages over UDP.

- The ASA does not have the ability to recognize HTTP cloaking where RTSP messages are hidden in the HTTP messages.
- With Cisco IP/TV, the number of translates the ASA performs on the SDP part of the message is proportional to the number of program listings in the Content Manager (each program listing can have at least six embedded IP addresses).
- You can configure NAT for Apple QuickTime 4 or RealPlayer. Cisco IP/TV only works with NAT if the Viewer and Content Manager are on the outside network and the server is on the inside network.

Configure RTSP Inspection Policy Map

You can create an RTSP inspection policy map to customize RTSP inspection actions if the default inspection behavior is not sufficient for your network.

You can optionally create a RTSP inspection class map to define the traffic class for RTSP inspection. The other option is to define the traffic classes directly in the RTSP inspection policy map. The difference between creating a class map and defining the traffic match directly in the inspection map is that you can create more complex match criteria and you can reuse class maps. Although this procedure explains inspection maps, the matching criteria used in class maps are the same as those explained in the step relating to the **Inspection** tab. You can configure RTSP class maps by selecting **Configuration** > **Firewall** > **Objects** > **Class Maps** > **RTSP**, or by creating them while configuring the inspection map.

Tip You can configure inspection maps while creating service policies, in addition to the procedure explained below. The contents of the map are the same regardless of how you create it.

Before you begin

Some traffic matching options use regular expressions for matching purposes. If you intend to use one of those techniques, first create the regular expression or regular expression class map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **RTSP**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map to and click **Edit**.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 Click the **Parameters** tab and configure the desired options:

- **Enforce Reserve Port Protection**—Whether to restrict the use of reserved ports during media port negotiation.
- **Maximum URL Length**—The maximum length of the URL allowed in the message, 0 to 6000.

Step 5 Click the **Inspections** tab and define the specific inspections you want to implement based on traffic characteristics.

You can define traffic matching criteria based on RTSP class maps, by configuring matches directly in the inspection map, or both.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose **Single Match** to define the criterion directly, or **Multiple Match**, in which case you select the RTSP class map that defines the criteria.

c) If you are defining the criterion here, choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map. Then, configure the criterion as follows:

- **URL Filter**—Match the URL against the selected regular expression or regular expression class.
- **Request Method**—Match the request method: announce, describe, get_parameter, options, pause, play, record, redirect, setup, set_parameters, teardown.

d) Choose the action to take for matching traffic. For URL matching, you can drop the connection or log it, and you can enable logging of dropped connections. For Request Method matches, you can apply a rate limit in packets per second.

e) Click **OK** to add the inspection. Repeat the process as needed.

Step 6 Click **OK** in the RTSP Inspect Map dialog box.

You can now use the inspection map in an RTSP inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

SIP Inspection

SIP is a widely used protocol for Internet conferencing, telephony, presence, events notification, and instant messaging. Partially because of its text-based nature and partially because of its flexibility, SIP networks are subject to a large number of security threats.

SIP application inspection provides address translation in message header and body, dynamic opening of ports and basic sanity checks. It also supports application security and protocol conformance, which enforce the sanity of the SIP messages, as well as detect SIP-based attacks.

SIP inspection is enabled by default. You need to configure it only if you want non-default processing, or if you want to identify a TLS proxy to enable encrypted traffic inspection. The following topics explain SIP inspection in more detail.

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 signaling. SDP specifies the ports for the media stream. Using SIP, the ASA can support any SIP VoIP gateways and VoIP proxy servers. SIP and SDP are defined in the following RFCs:

- SIP: Session Initiation Protocol, RFC 3261
- SDP: Session Description Protocol, RFC 2327

To support SIP calls through the ASA, signaling messages for the media connection addresses, media ports, and embryonic connections for the media must be inspected, because while the signaling is sent over a well-known destination port (UDP/TCP 5060), the media streams are dynamically allocated. Also, SIP embeds IP addresses in the user-data portion of the IP packet. Note that the maximum length of the SIP Request URI that the ASA supports is 255.

Instant Messaging (IM) applications also use SIP extensions (defined in RFC 3428) and SIP-specific event notifications (RFC 3265). After users initiate a chat session (registration/subscription), the IM applications use the MESSAGE/INFO methods and 202 Accept responses when users chat with each other. 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 SIP inspection supports the Chat feature only. Whiteboard, File Transfer, and Application Sharing are not supported. RTC Client 5.0 is not supported.

Limitations for SIP Inspection

SIP inspection is tested and supported for Cisco Unified Communications Manager (CUCM) 7.0, 8.0, 8.6, and 10.5. It is not supported for CUCM 8.5, or 9.x. SIP inspection might work with other releases and products.

SIP inspection does not support the T.38 MIME Internet Facsimile Protocol(IFP). SIP inspection drops SIP invitations that use the T.38 MIME audio sub-type. If you need to allow this type, disable SIP inspection and write an access control rule that allows the RTP streams.

NAT Limitations for SIP Inspection

- SIP inspection applies NAT for embedded IP addresses. However, if you configure NAT to translate both source and destination addresses, the external address ("from" in the SIP header for the "trying" response message) is not rewritten. Thus, you should use object NAT when working with SIP traffic so that you avoid translating the destination address.
- Do not configure NAT or PAT for interfaces with the same, or lower (source) to higher (destination), security levels. This configuration is not supported.

- If you configure SIP inspection for a targeted traffic class (that is, not the inspection_default traffic class), ensure that you use a bi-directional ACL and that you specify the 5060 destination port only. Otherwise, you might have NAT problems where the IP address in the SIP header is not translated even though the IP packet is correctly translated.
- If you hard-code the mapped address in the VIA header of the SIP invitation, do not enable SIP inspection. You can run into problems if you use static NAT to translate the source client address, and the interface for the default route is different than the interface for the connected route used by the client.

PAT Limitations for SIP Inspection

The following limitations and restrictions apply when using PAT with SIP:

- If a remote endpoint tries to register with a SIP proxy on a network protected by the ASA, the registration fails under very specific conditions, as follows:
  - PAT is configured for the remote endpoint.
  - The SIP registrar server is on the outside network.
  - The port is missing in the contact field in the REGISTER message sent by the endpoint to the proxy server.
- If a SIP device transmits a packet in which the SDP portion has an IP address in the owner/creator field (o=) that is different than the IP address in the connection field (c=), the IP address in the o= field may not be properly translated. This is due to a limitation in the SIP protocol, which does not provide a port value in the o= field. Because PAT needs a port to translate, the translation fails.
- When using PAT, any SIP header field which contains an internal IP address without a port might not be translated and hence the internal IP address will be leaked outside. If you want to avoid this leakage, configure NAT instead of PAT.

Default SIP Inspection

SIP inspection is enabled by default using the default inspection map, which includes the following:

- SIP instant messaging (IM) extensions: Enabled.
- Non-SIP traffic on SIP port: Dropped.
- Hide server's and endpoint's IP addresses: Disabled.
- Mask software version and non-SIP URIs: Disabled.
- Ensure that the number of hops to destination is greater than 0: Enabled.
- RTP conformance: Not enforced.
- SIP conformance: Do not perform state checking and header validation.

Also note that inspection of encrypted traffic is not enabled. You must configure a TLS proxy to inspect encrypted traffic.

Configure SIP Inspection Policy Map

You can create a SIP inspection policy map to customize SIP inspection actions if the default inspection behavior is not sufficient for your network.

You can optionally create a SIP inspection class map to define the traffic class for SIP inspection. The other option is to define the traffic classes directly in the SIP inspection policy map. The difference between creating a class map and defining the traffic match directly in the inspection map is that you can create more complex match criteria and you can reuse class maps. Although this procedure explains inspection maps, the matching criteria used in class maps are the same as those explained in the step relating to the **Inspection** tab. You can configure SIP class maps by selecting **Configuration** > **Firewall** > **Objects** > **Class Maps** > **SIP**, or by creating them while configuring the inspection map.

Tip You can configure inspection maps while creating service policies, in addition to the procedure explained below. The contents of the map are the same regardless of how you create it.

Before you begin

Some traffic matching options use regular expressions for matching purposes. If you intend to use one of those techniques, first create the regular expression or regular expression class map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **SIP**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map to view its contents. You can change the security level directly, or click **Customize** to edit the map. The remainder of the procedure assumes you are customizing or adding a map.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 In the **Security Level** view of the SIP Inspect Map dialog box, select the level that best matches your desired configuration. The default level is Low.

If one of the preset levels matches your requirements, you are now done. Just click **OK**, skip the rest of this procedure, and use the map in a service policy rule for SIP inspection.

Step 5 If you need to customize the settings further, click **Details**, and do the following:

a) Click the **Filtering** tab and choose whether to enable SIP instant messaging (IM) extensions or to permit non-SIP traffic on the SIP port.
b) Click the **IP Address Privacy** tab and choose whether to hide the server and endpoint IP addresses.
c) Click the **Hop Count** tab and choose whether to ensure that the number of hops to the destination is greater than 0. This checks the value of the Max-Forwards header, which cannot be zero before reaching the destination. You must also choose the action to take for non-conforming traffic (drop packet, drop connection, reset, or log) and whether to enable or disable logging.

d) Click the **RTP Conformance** tab and choose whether to check RTP packets that are flowing on the pinholes for protocol conformance. If you check for conformance, you can also choose whether to limit the payload to audio or video, based on the signaling exchange.

e) Click the **SIP Conformance** tab and choose whether to enable state transition checking and strict validation of header fields. For each option you choose, select the action to take for non-conforming traffic (drop packet, drop connection, reset, or log) and whether to enable or disable logging.

f) Click the **Field Masking** tab and choose whether to inspect non-SIP URIs in Alert-Info and Call-Info headers and to inspect the server's and endpoint's software version in the User-Agent and Server headers. For each option you choose, select the action to take (mask or log) and whether to enable or disable logging.

g) Click the **TVS Server** tab and identify the Trust Verification Services servers, which enable Cisco Unified IP Phones to authenticate application servers during HTTPS establishment. You can identify up to four servers; enter their IP addresses separated by commas. SIP inspection opens pinholes to each server for each registered phone, and the phone decides which to use.

   Configure the Trust Verification Services server on the CUCM server. If the configuration uses a non-default port, enter the port number (in the range 1026 to 32768). The default port is 2445.

Step 6 Click the **Inspections** tab and define the specific inspections you want to implement based on traffic characteristics.

You can define traffic matching criteria based on SIP class maps, by configuring matches directly in the inspection map, or both.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose **Single Match** to define the criterion directly, or **Multiple Match**, in which case you select the SIP class map that defines the criteria.

c) If you are defining the criterion here, choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map. Then, configure the criterion as follows:

- **Called Party**—Match the called party, as specified in the To header, against the selected regular expression or regular expression class.
- **Calling Party**—Match the calling party, as specified in the From header, against the selected regular expression or regular expression class.
- **Content Length**—Match a SIP content header of a length greater than specified, between 0 and 65536 bytes.
- **Content Type**—Match the Content Type header, either the SDP type or a type that matches the selected regular expression or regular expression class.
- **IM Subscriber**—Match the SIP IM subscriber against the selected regular expression or regular expression class.
- **Message Path**—Match the SIP Via header against the selected regular expression or regular expression class.

- **Request Method**—Match the SIP request method: ack, bye, cancel, info, invite, message, notify, options, prack, refer, register, subscribe, unknown, update.
- **Third-Party Registration**—Match the requester of a third-party registration against the selected regular expression or regular expression class.
- **URI Length**—Match a URI in the SIP headers of the selected type (SIP or TEL) that is greater than the length specified, between 0 and 65536 bytes.

d) Choose the action to take for matching traffic (drop packet, drop connection, reset, log) and whether to enable or disable logging. For Request Method matches to “invite” and “register,” you can also apply a rate limit in packets per second.
e) Click **OK** to add the inspection. Repeat the process as needed.

Step 7 Click **OK** in the SIP Inspect Map dialog box.

You can now use the inspection map in a SIP inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

Skinny (SCCP) Inspection

SCCP (Skinny) application inspection performs translation of embedded IP address and port numbers within the packet data, and dynamic opening of pinholes. It also performs additional protocol conformance checks and basic state tracking.

SCCP inspection is enabled by default. You need to configure it only if you want non-default processing, or if you want to identify a TLS proxy to enable encrypted traffic inspection.

The following sections describe SCCP application inspection.

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.

The ASA supports PAT and NAT for SCCP. PAT is necessary if you have more IP phones than global IP addresses for the IP phones to use. By supporting NAT and PAT of SCCP Signaling packets, Skinny application inspection ensures that all SCCP signaling and media packets can traverse the ASA.

Normal traffic between Cisco CallManager and Cisco IP Phones uses SCCP and is handled by SCCP inspection without any special configuration. The ASA also supports DHCP options 150 and 66, which it accomplishes by sending the location of a TFTP server to Cisco IP Phones and other DHCP clients. Cisco IP Phones might also include DHCP option 3 in their requests, which sets the default route.

Note The ASA supports inspection of traffic from Cisco IP Phones running SCCP protocol version 22 and earlier.

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. A 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 ACL 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 ACL or static entry is required to allow the Cisco IP Phones to initiate the connection.

Limitations for SCCP Inspection

SCCP inspection is tested and supported for Cisco Unified Communications Manager (CUCM) 7.0, 8.0, 8.6, and 10.5. It is not supported for CUCM 8.5, or 9.x. SCCP inspection might work with other releases and products.

If the address of an internal Cisco CallManager is configured for NAT or PAT to a different IP address or port, registrations for external Cisco IP Phones fail because the ASA does not support NAT or PAT for the file content transferred over TFTP. Although the ASA supports NAT of TFTP messages and opens a pinhole for the TFTP file, the ASA cannot translate the Cisco CallManager IP address and port embedded in the Cisco IP Phone configuration files that are transferred by TFTP during phone registration.

Note The ASA supports stateful failover of SCCP calls except for calls that are in the middle of call setup.

Default SCCP Inspection

SCCP inspection is enabled by default using these defaults:

- Registration: Not enforced.
- Maximum message ID: 0x181.
- Minimum prefix length: 4
- Media timeout: 00:05:00
- Signaling timeout: 01:00:00.

- RTP conformance: Not enforced.

Configure a Skinny (SCCP) Inspection Policy Map

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.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **SCCP (Skinny)**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map to view its contents. You can change the security level directly, or click **Customize** to edit the map. The remainder of the procedure assumes you are customizing or adding a map.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 In the **Security Level** view of the SCCP (Skinny) Inspect Map dialog box, select the level that best matches your desired configuration. The default level is Low.

If one of the preset levels matches your requirements, you are now done. Just click **OK**, skip the rest of this procedure, and use the map in a service policy rule for SCCP inspection.

Step 5 If you need to customize the settings further, click **Details**, and do the following:

a) Click the **Parameters** tab and choose among the following options:

- **Enforce endpoint registration**—Whether Skinny endpoints must register before placing or receiving calls.
- **Maximum Message ID**—The maximum SCCP station message ID allowed. The default maximum is 0x181. The hex number can be 0x0 to 0xffff.
- **SCCP Prefix Length**—The maximum and minimum SCCP prefix length. The default minimum is 4; there is no default maximum.
- **Timeouts**—Whether to set timeouts for media and signaling connections, and the value of those timeouts. The defaults are 5 minutes for media, 1 hour for signaling.

b) Click the **RTP Conformance** tab and choose whether to check RTP packets that are flowing on the pinholes for protocol conformance. If you check for conformance, you can also choose whether to limit the payload to audio or video, based on the signaling exchange.

Step 6 (Optional) Click the **Message ID Filtering** tab to identify traffic to drop based on the station message ID field in SCCP messages.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion).
c) In the **Value** fields, identify the traffic based on the station message ID value in hexadecimal, from 0x0 to 0xffff. Either enter the value for a single message ID, or enter the beginning and ending value for a range of IDs.
d) Choose whether to enable or disable logging. The action is always to drop the packet.
e) Click **OK** to add the filter. Repeat the process as needed.

Step 7 Click **OK** in the SCCP (Skinny) Inspect Map dialog box.

You can now use the inspection map in an SCCP inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure Application Layer Protocol Inspection, on page 306.

STUN Inspection

Session Traversal Utilities for NAT (STUN), defined in RFC 5389, is used by WebRTC clients for browser-based real-time communications so that plug-ins are not necessary. WebRTC clients often use cloud STUN servers to learn their public IP addresses and ports. WebRTC uses Interactive Connectivity Establishment (ICE, RFC 5245) to verify connectivity between clients. These clients typically use UDP, although they can also use TCP or other protocols.

Because firewalls often block outgoing UDP traffic, WebRTC products such as Cisco Spark can have problems completing connections. STUN inspection opens pinholes for STUN endpoints, and enforces basic STUN and ICE compliance, to allow communications for clients if the connectivity check is acknowledged by both sides. Thus, you can avoid opening new ports in your access rules to enable these applications.

When you enable STUN inspection on the default inspection class, TCP/UDP port 3478 is watched for STUN traffic. The inspection supports IPv4 addresses and TCP/UDP only.

There are some NAT limitations for STUN inspection. For WebRTC traffic, static NAT/PAT44 are supported. Cisco Spark can support additional types of NAT, because Spark does not require pinholes. You can also use NAT/PAT64, including dynamic NAT/PAT, with Cisco Spark.

STUN inspection is supported in failover and cluster modes, as pinholes are replicated. However, the transaction ID is not replicated among nodes. In the case where a node fails after receiving a STUN Request and another node received the STUN Response, the STUN Response will be dropped.



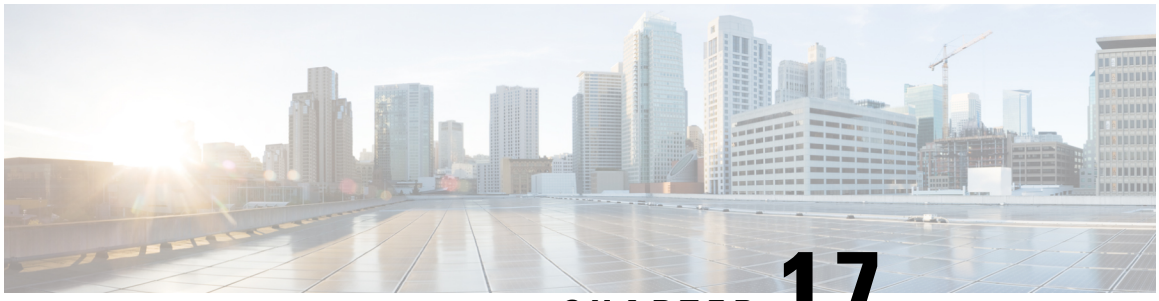
Note STUN inspection uses transaction IDs to match requests and responses. If you use debug to troubleshoot connection drops, note that the system changes the format (endianness) of the IDs for the debug output, so they do not compare directly to those you might see in a pcap.

For information on enabling STUN inspection, see Configure Application Layer Protocol Inspection, on page 306.

History for Voice and Video Protocol Inspection

Feature Name	Releases	Feature Information
SIP, SCCP, and TLS Proxy support for IPv6	9.3(1)	You can now inspect IPv6 traffic when using SIP, SCCP, and TLS Proxy (using SIP or SCCP). We did not modify any ASDM screens.
SIP support for Trust Verification Services, NAT66, CUCM 10.5, and model 8831 phones.	9.3(2)	You can now configure Trust Verification Services servers in SIP inspection. You can also use NAT66. SIP inspection has been tested with CUCM 10.5. We added Trust Verification Services Server support to the SIP inspection policy map.
Improved SIP inspection performance on multiple core ASA.	9.4(1)	If you have multiple SIP signaling flows going through an ASA with multiple cores, SIP inspection performance has been improved. However, you will not see improved performance if you are using a TLS, phone, or IME proxy. We did not modify any ASDM screens.
SIP inspection support in ASA clustering	9.4(1)	You can now configure SIP inspection on the ASA cluster. A control flow can be created on any unit (due to load balancing), but its child data flows must reside on the same unit. TLS Proxy configuration is not supported. We did not modify any screens.
SIP inspection support for Phone Proxy and UC-IME Proxy was removed.	9.4(1)	You can no longer use Phone Proxy or UC-IME Proxy when configuring SIP inspection. Use TLS Proxy to inspect encrypted traffic. We removed Phone Proxy and UC-IME Proxy from the Select SIP Inspect Map service policy dialog box.
H.323 inspection support for the H.255 FACILITY message coming before the H.225 SETUP message for H.460.18 compatibility.	9.6(1)	You can now configure an H.323 inspection policy map to allow for H.225 FACILITY messages to come before the H.225 SETUP message, which can happen when endpoints comply with H.460.18. We added an option to the Call Attributes tab in the H.323 inspection policy map.
Session Traversal Utilities for NAT (STUN) inspection.	9.6(2)	You can now inspect STUN traffic for WebRTC applications including Cisco Spark. Inspection opens pinholes required for return traffic. We added an option to the Rule Actions > Protocol Inspection tab of the Add/Edit Service Policy dialog box.

Feature Name	Releases	Feature Information
Support for TLSv1.2 in TLS proxy and Cisco Unified Communications Manager 10.5.2.	9.7(1)	You can now use TLSv1.2 with TLS proxy for encrypted SIP or SCCP inspection with the Cisco Unified Communications Manager 10.5.2. The TLS proxy supports the additional TLSv1.2 cipher suites added as part of the **client cipher-suite** command. We did not modify any screens.
TLS proxy deprecated for SCCP (Skinny) inspection.	9.13(1)	The **tls-proxy** keyword, and support for SCCP/Skinny encrypted inspection, was deprecated. The keyword will be removed from the **inspect skinny** command in a future release.
TLS proxy support eliminated for SCCP (Skinny) inspection.	9.14(1)	The **tls-proxy** keyword, and support for SCCP/Skinny encrypted inspection, was removed.
The default SIP inspection policy map drops non-SIP traffic.	9.16(1)	For SIP-inspected traffic, the default is now to drop non-SIP traffic. The previous default was to allow non-SIP traffic on ports inspected for SIP. We changed the default SIP policy map to include the **no traffic-non-sip** command.



CHAPTER 17

Inspection for Mobile Networks

The following topics explain application inspection for protocols used in mobile networks such as LTE. These inspections require the Carrier license. For information on why you need to use inspection for certain protocols, and the overall methods for applying inspection, see Getting Started with Application Layer Protocol Inspection, on page 299.

- Mobile Network Inspection Overview, on page 373
- Licensing for Mobile Network Protocol Inspection, on page 380
- Defaults for GTP Inspection, on page 380
- Configure Mobile Network Inspection, on page 381
- Monitoring Mobile Network Inspection, on page 402
- History for Mobile Network Inspection, on page 406

Mobile Network Inspection Overview

The following topics explain the inspections available for protocols used in mobile networks such as LTE. There are other services available for SCTP traffic in addition to inspection.

GTP Inspection Overview

GPRS Tunneling Protocol is used in GSM, UMTS and LTE networks for general packet radio service (GPRS) traffic. GTP provides a tunnel control and management protocol to provide GPRS network access for a mobile station by creating, modifying, and deleting tunnels. GTP also uses a tunneling mechanism for carrying user data packets.

Service provider networks use GTP to tunnel multi-protocol packets through the GPRS backbone between endpoints. In GTPv0-1, GTP is used for signaling between gateway GPRS support nodes (GGSN) and serving GPRS support nodes (SGSN). In GTPv2, the signaling is between Packet Data Network Gateways (PGW) and the Serving Gateway (SGW) as well as other endpoints. The GGSN/PGW is the interface between the GPRS wireless data network and other networks. The SGSN/SGW performs mobility, data session management, and data compression.

You can use the ASA to provide protection against rogue roaming partners. Place the device between the home GGSN/PGW and visited SGSN/SGW endpoints and use GTP inspection on the traffic. GTP inspection works only on traffic between these endpoints. In GTPv2, this is known as the S5/S8 interface.

GTP and associated standards are defined by 3GPP (3rd Generation Partnership Project). For detailed information, see http://www.3gpp.org.

Tracking Location Changes for Mobile Stations

You can use GTP inspection to track location changes for mobile stations. Tracking location changes might help you identify fraudulent roaming charges, for example, if you see a mobile station move from one location to another within an unlikely time window, such as moving from a cell in the United States to one in Europe within 30 minutes.

When you enable location logging, the system generates syslog messages for new or changed location for each International Mobile Subscriber Identity (IMSI):

- 324010 indicates the creation of a new PDP context, and includes the Mobile Country Code (MCC), Mobile Network Code (MNC), the information elements, and optionally the cell ID where the user currently is registered. The cell ID is extracted from the Cell Global Identification (CGI) or E-UTRAN Cell Global Identifier (ECGI).
- 324011 indicates that the IMSI has moved from the one stored during the PDP context creation. The message shows the previous and current MCC/MNC, information elements, and optionally, cell ID.

By default, syslog messages do not include timestamp information. If you plan to analyze these messages to identify improbable roaming, you must also enable timestamps. Timestamp logging is not part of the GTP inspection map. Go to **Configuration** > **Device Management** > **Logging** > **Syslog Setup** and select the **Enable Timestamp on Syslog Messages** option.

For information on enabling location logging, see Configure a GTP Inspection Policy Map, on page 381.

GTP Inspection Limitations

Following are some limitations on GTP inspection:

- GTPv2 piggybacking messages are not supported. They are always dropped.
- GTPv2 emergency UE attach is supported only if it contains IMSI (International Mobile Subscriber Identity).
- GTP inspection does not inspect early data. That is, data sent from a PGW or SGW right after a Create Session Request but before the Create Session Response.
- For GTPv2, inspection supports up to 3GPP 29.274 V15.5.0. For GTPv1, support is up to 3GPP 29.060 V15.2.0. For GTPv0, support is up to release 8.
- GTP inspection does not support inter-SGSN handoff to the secondary PDP context. Inspection needs to do the handoff for both primary and secondary PDP contexts.
- When you enable GTP inspection, connections that use GTP-in-GTP encapsulation are always dropped.

Stream Control Transmission Protocol (SCTP) Inspection and Access Control

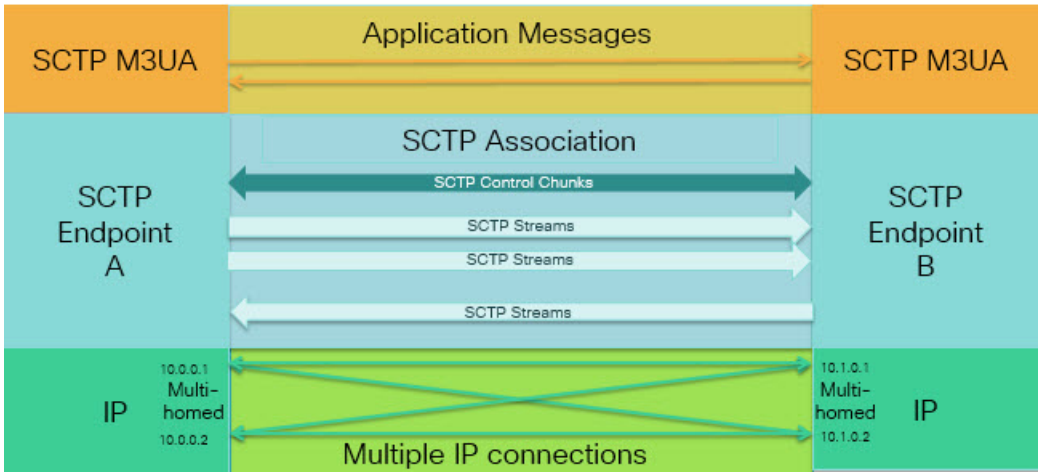
SCTP (Stream Control Transmission Protocol) is described in RFC 4960. The protocol supports the telephony signaling protocol SS7 over IP and is also a transport protocol for several interfaces in the 4G LTE mobile network architecture.

SCTP is a transport-layer protocol operating on top of IP in the protocol stack, similar to TCP and UDP. However, SCTP creates a logical communication channel, called an association, between two end nodes over one or more source or destination IP addresses. This is called multi-homing. An association defines a set of IP addresses on each node (source and destination) and a port on each node. Any IP address in the set can be

used as either a source or a destination IP address of data packets associated to this association to form multiple connections. Within each connection, multiple streams may exist to send messages. A stream in SCTP represents a logical application data channel.

The following figure illustrates the relationship between an association and its streams.

Figure 48: Relationship Between SCTP Association and Streams



If you have SCTP traffic going through the ASA, you can control access based on SCTP ports, and implement application layer inspection to enable connections and to optionally filter on payload protocol ID to selectively drop, log, or rate limit applications.

Note Each node can have up to three IP addresses. Any addresses over the limit of three are ignored and not included in the association. Pinholes for secondary IP addresses are opened automatically. You do not need to write access control rules to allow them.

Heartbeats and Association State

To maintain the association and ensure accurate reporting, the inspection engine monitors SCTP heartbeats. The system links the connection duration to the heartbeat signals. If the heartbeat is lost or the association is explicitly terminated, the device closes the connection and finalizes the event duration.

SCTP Association and Port Reporting

SCTP inspection is association-based rather than flow-based. Because the device tracks the entire association as a single unit, it does not map individual stream ports to traditional port fields in the logging output. As a result, when viewing connection details, you may observe the source and destination ports displayed as `0`. This is because the association is uniquely identified by the IP addresses and the association ID, rather than by individual source/destination port pairs.

SCTP Stateful Inspection

Similar to TCP, SCTP traffic is automatically inspected at layer 4 to ensure well-structured traffic and limited RFC 4960 enforcement. The following protocol elements are inspected and enforced:

- Chunk types, flags, and length.

- Verification tags.
- Source and destination ports, to prevent association redirect attacks.
- IP addresses.

SCTP stateful inspection accepts or rejects packets based on the association state:

- Validating the 4-way open and close sequences for initial association establishment.
- Verifying the forward progression of TSN within an association and a stream.
- Terminating an association when seeing the ABORT chunk due to heartbeat failure. SCTP endpoints might send the ABORT chunk in response to bombing attacks.

If you decide you do not want these enforcement checks, you can configure SCTP state bypass for specific traffic classes, as explained in Configure Connection Settings for Specific Traffic Classes (All Services), on page 428.

SCTP Access Control

You can create access rules for SCTP traffic. These rules are similar to TCP/UDP port-based rules, where you simply use **sctp** as the protocol, and the port numbers are SCTP ports. You can create service objects or groups for SCTP, or specify the ports directly. See the following topics.

- Configure Service Objects and Service Groups, on page 33
- Configure Extended ACLs, on page 47
- Configure Access Rules, on page 19

SCTP NAT

You can apply static network object NAT to the addresses in SCTP association establishment messages. Although you can configure static twice NAT, this is not recommended because the topology of the destination part of the SCTP association is unknown. You cannot use dynamic NAT/PAT.

NAT for SCTP depends upon SCTP stateful inspection rather than SCTP application-layer inspection. Thus, you cannot NAT traffic if you configure SCTP state bypass.

SCTP Application Layer Inspection

You can further refine your access rules by enabling SCTP inspection and filtering on SCTP applications. You can selectively drop, log, or rate limit SCTP traffic classes based on the payload protocol identifier (PPID).

If you decide to filter on PPID, keep the following in mind:

- PPIDs are in data chunks, and a given packet can have multiple data chunks or even a control chunk. If a packet includes a control chunk or multiple data chunks, the packet will not be dropped even if the assigned action is drop.
- If you use PPID filtering to drop or rate-limit packets, be aware that the transmitter will resend any dropped packets. Although a packet for a rate-limited PPID might make it through on the next attempt, a packet for a dropped PPID will again be dropped. You might want to evaluate the eventual consequence of these repeated drops on your network.

SCTP Limitations

SCTP support includes the following limitations.

- Each node can have up to three IP addresses. Any addresses over the limit of three are ignored and not included in the association. Pinholes for secondary IP addresses are opened automatically. You do not need to write access control rules to allow them.
- Unused pinholes time out in 5 minutes.
- Dual stack IPv4 and IPv6 addresses on multi-homed endpoints is not supported.
- Network object static NAT is the only supported type of NAT. Also, NAT46 and NAT64 are not supported.
- Fragmentation and reassembly of SCTP packets is done only for traffic handled by Diameter, M3UA, and SCTP PPID-based inspection.
- ASCONF chunks, which are used to dynamically add or delete IP addresses in SCTP, are not supported.
- The Hostname parameter in INIT and INIT-ACK SCTP messages, which is used to specify a hostname which can then be resolved to an IP address, is not supported.
- SCTP/M3UA does not support equal-cost multipath routing (ECMP), whether configured on the ASA or elsewhere in the network. With ECMP, packets can be routed to a destination over multiple best paths. However, an SCTP/M3UA packet response to a single destination has to come back on the same interface that it exited. Even though the response can come from any M3UA server, it must always come back on the same interface that it exited. The symptom for this problem is that SCTP INIT-ACK packets are dropped, which you can see in the **show asp drop flow sctp-chunk-init-timeout** counter:

  ```
  Flow drop:
  SCTP INIT timed out (not receiving INIT ACK)(sctp-chunk-init-timeout)
  ```

  If you encounter this problem, you can resolve it by configuring static routes to the M3UA servers, or by configuring policy-based routing to implement a network design that ensures that INIT-ACK packets go through the same interface as the INIT packets.

Diameter Inspection

Diameter is an Authentication, Authorization, and Accounting (AAA) protocol used in next-generation mobile and fixed telecom networks such as EPS (Evolved Packet System) for LTE (Long Term Evolution) and IMS (IP Multimedia Subsystem). It replaces RADIUS and TACACS in these networks.

Diameter uses TCP and SCTP as the transport layer, and secures communications using TCP/TLS and SCTP/DTLS. It can optionally provide data object encryption as well. For detailed information on Diameter, see RFC 6733.

Diameter applications perform service management tasks such as deciding user access, service authorization, quality of service, and rate of charging. Although Diameter applications can appear on many different control-plane interfaces in the LTE architecture, the ASA inspects Diameter command codes and attribute-value pairs (AVP) for the following interfaces only:

- S6a: Mobility Management Entity (MME) - Home Subscription Service (HSS).
- S9: PDN Gateway (PDG) - 3GPP AAA Proxy/Server.
- Rx: Policy Charging Rules Function (PCRF) - Call Session Control Function (CSCF).

Diameter inspection opens pinholes for Diameter endpoints to allow communication. The inspection supports 3GPP version 12 and is RFC 6733 compliant. You can use it for TCP/TLS (by specifying a TLS proxy when you enable inspection) and SCTP, but not SCTP/DTLS. Use IPsec to provide security to SCTP Diameter sessions.

You can optionally use a Diameter inspection policy map to filter traffic based on application ID, command codes, and AVP, to apply special actions such as dropping packets or connections, or logging them. You can create custom AVP for newly-registered Diameter applications. Filtering lets you fine-tune the traffic you allow on your network.

Note Diameter messages for applications that run on other interfaces will be allowed and passed through by default. However, you can configure a Diameter inspection policy map to drop these applications by application ID, although you cannot specify actions based on the command codes or AVP for these unsupported applications.

M3UA Inspection

MTP3 User Adaptation (M3UA) is a client/server protocol that provides a gateway to the SS7 network for IP-based applications that interface with the SS7 Message Transfer Part 3 (MTP3) layer. M3UA makes it possible to run the SS7 User Parts (such as ISUP) over an IP network. M3UA is defined in RFC 4666.

M3UA uses SCTP as the transport layer. SCTP port 2905 is the default port.

The MTP3 layer provides networking functions such as routing and node addressing, but uses point codes to identify nodes. The M3UA layer exchanges Originating Point Codes (OPC) and Destination Point Codes (DPC). This is similar to how IP uses IP addresses to identify nodes.

M3UA inspection provides limited protocol conformance. You can optionally implement strict application server process (ASP) state checking and additional message validation for select messages. Strict ASP state checking is required if you want stateful failover or if you want to operate within a cluster. However, strict ASP state checking works in Override mode only, it does not work if you are running in Loadsharing or Broadcast mode (per RFC 4666). The inspection assumes there is one and only one ASP per endpoint.

You can optionally apply access policy based on point codes or Service Indicators (SI). You can also apply rate limiting based on message class and type.

M3UA Protocol Conformance

M3UA inspection provides the following limited protocol enforcement. Inspection drops and logs packets that do not meet requirements.

- Common message header. Inspection validates all fields in the common header.
  - Version 1 only.
  - Message length must be correct.
  - Message type class with a reserved value is not allowed.
  - Invalid message ID within the message class is not allowed.
- Payload data message.
  - Only one parameter of a given type is allowed.

- Data messages on SCTP stream 0 are not allowed.

- The Affected Point Code field must be present in the following messages or the message is dropped: Destination Available (DAVA), Destination Unavailable (DUNA), Destination State Audit (DAUD), Signaling Congestion (SCON), Destination User Part Unavailable (DUPU), Destination Restricted (DRST).
- If you enable message tag validation for the following messages, the content of certain fields are checked and validated. Messages that fail validation are dropped.
  - Destination User Part Unavailable (DUPU)—The User/Cause field must be present, and it must contain only valid cause and user codes.
  - Error—All mandatory fields must be present and contain only allowed values. Each error message must contain the required fields for that error code.
  - Notify—The status type and status information fields must contain allowed values only.
- If you enable strict application server process (ASP) state validation, the system maintains the ASP states of M3UA sessions and allows or drops ASP messages based on the validation result. If you do not enable strict ASP state validation, all ASP messages are forwarded uninspected.

M3UA Inspection Limitations

Following are some limitations on M3UA inspection.

- NAT is not supported for IP addresses that are embedded in M3UA data.
- M3UA strict application server process (ASP) state validation depends on SCTP stateful inspection. Do not implement SCTP state bypass and M3UA strict ASP validation on the same traffic.
- Strict ASP state checking is required if you want stateful failover or if you want to operate within a cluster. However, strict ASP state checking works in Override mode only, it does not work if you are running in Loadsharing or Broadcast mode (per RFC 4666). The inspection assumes there is one and only one ASP per endpoint.

RADIUS Accounting Inspection Overview

The purpose of RADIUS accounting inspection is to prevent over-billing attacks on GPRS networks that use RADIUS servers. Although you do not need the Carrier license to implement RADIUS accounting inspection, it has no purpose unless you are implementing GTP inspection and you have a GPRS setup.

The over-billing attack in GPRS networks results in consumers being billed for services that they have not used. In this case, a malicious attacker sets up a connection to a server and obtains an IP address from the SGSN. When the attacker ends the call, the malicious server will still send packets to it, which gets dropped by the GGSN, but the connection from the server remains active. The IP address assigned to the malicious attacker gets released and reassigned to a legitimate user who will then get billed for services that the attacker will use.

RADIUS accounting inspection prevents this type of attack by ensuring the traffic seen by the GGSN is legitimate. With the RADIUS accounting feature properly configured, the ASA 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 ASA 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 ASA can validate the message. If the shared secret is not configured, the ASA will only check that the source IP address is one of the configured addresses allowed to send the RADIUS messages.



Note When using RADIUS accounting inspection with GPRS enabled, the ASA checks for the 3GPP-Session-Stop-Indicator in the Accounting Request STOP messages to properly handle secondary PDP contexts. Specifically, the ASA requires that the Accounting Request STOP messages include the 3GPP-SGSN-Address attribute before it will terminate the user sessions and all associated connections. Some third-party GGSNs might not send this attribute by default.

Licensing for Mobile Network Protocol Inspection

Inspection of the following protocols requires the license listed in the table below.

- GTP
- SCTP.
- Diameter
- M3UA

Model	License Requirement
ASAv (all models)	Carrier license (enabled by default)
Firepower 4100	Carrier license
Firepower 9300	Carrier license
All other models	The Carrier license is not available on other models. You cannot inspect these protocols.

Defaults for GTP Inspection

GTP inspection is not enabled by default. However, if you enable it without specifying your own inspection map, a default map is used that provides the following processing. You need to configure a map only if you want different values.

- Errors are not permitted.
- The maximum number of requests is 200.
- The maximum number of tunnels is 500. This is equivalent to the number of PDP contexts (endpoints).
- The GTP endpoint timeout is 30 minutes. Endpoints include GSNs (GTPv0,1) and SGW/PGW (GTPv2).

- The PDP context timeout is 30 minutes. In GTPv2, this is the bearer context timeout.
- The request timeout is 1 minute.
- The signaling timeout is 30 minutes.
- The tunneling timeout is 1 hour.
- The T3 response timeout is 20 seconds.
- Unknown message IDs are allowed. You can configure **match message v1/v2 id range** commands to drop and log any commands that you do not support or want to allow. Messages are considered unknown if they are either undefined or are defined in GTP releases that the system does not support.

Configure Mobile Network Inspection

Inspections for protocols used in mobile networks are not enabled by default. You must configure them if you want to support mobile networks.

Procedure

Step 1 (Optional.) Configure a GTP Inspection Policy Map, on page 381.

Step 2 (Optional.) Configure an SCTP Inspection Policy Map, on page 385.

Step 3 (Optional.) Configure a Diameter Inspection Policy Map, on page 386.

If you want to filter on attribute-value pairs (AVP) that are not yet supported in the software, you can create custom AVP for use in the Diameter inspection policy map. See Create a Custom Diameter Attribute-Value Pair (AVP), on page 389.

Step 4 (Optional.) If you want to inspect encrypted Diameter TCP/TLS traffic, create the required TLS proxy as described in Inspecting Encrypted Diameter Sessions, on page 390

Step 5 (Optional.) Configure an M3UA Inspection Policy Map, on page 397

Step 6 Configure the Mobile Network Inspection Service Policy , on page 399.

Step 7 (Optional.) Configure RADIUS Accounting Inspection, on page 400.

RADIUS accounting inspection protects against over-billing attacks.

Configure a GTP Inspection Policy Map

If you want to enforce additional parameters on GTP traffic, and the default map does not meet your needs, create and configure a GTP map.

Before you begin

Some traffic matching options use regular expressions for matching purposes. If you intend to use one of those techniques, first create the regular expression or regular expression class map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **GTP**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map to view its contents. Click **Customize** to edit the map. The remainder of the procedure assumes you are customizing or adding a map.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 In the **Security Level** view of the GTP Inspect Map dialog box, view the current configuration of the map.

The view indicates whether the map uses default values or if you have customized it. If you need to customize the settings further, click **Details**, and continue with the procedure.

Tip
The **IMSI Prefix Filtering** button is a shortcut to configure IMSI prefix filtering, which is explained later in this procedure.

Step 5 Click the **Permit Parameters** tab and configure the desired options.

- **Permit Response**—When the ASA performs GTP inspection, by default the ASA drops GTP responses from GSNs or PGWs that were not specified in the GTP request. This situation occurs when you use load-balancing among a pool of GSN/PGW endpoints to provide efficiency and scalability of GPRS.

  To configure GSN/PGW pooling and thus support load balancing, create a network object group that specifies the GSN/PGW endpoints and select this as a "**From Object Group**." Likewise, create a network object group for the SGSN/SGW and select it as the "**To Object Group**." If the GSN/PGW responding belongs to the same object group as the GSN/PGW that the GTP request was sent to and if the SGSN/SGW is in an object group that the responding GSN/PGW is permitted to send a GTP response to, the ASA permits the response.

  The network object group can identify the endpoints by host address or by the subnet that contains them.

- **Permit Errors**—Whether to allows all packets, including invalid GTP packets or packets that otherwise would fail parsing and be dropped. Packets can still be dropped based on the actions you define in the policy map.

Step 6 Click the **General Parameters** tab and configure the desired options:

- **Maximum Number of Requests**—The maximum number of GTP requests that will be queued waiting for a response.
- **Maximum Number of Tunnels**—The maximum number of active GTP tunnels allowed. This is equivalent to the number of PDP contexts or endpoints. The default is 500. New requests will be dropped once the maximum number of tunnels is reached.
- **Enforce Timeout**—Whether to enforce idle timeouts for the following behaviors. Timeouts are in hh:mm:ss format.
  - Endpoint—The maximum period of inactivity before a GTP endpoint is removed.

- PDP-Context—The maximum period of inactivity before removing the PDP Context for a GTP session. In GTPv2, this is the bearer context.
- Request—The maximum period of inactivity after which a request is removed from the request queue. Any subsequent responses to a dropped request will also be dropped.
- Signaling—The maximum period of inactivity before GTP signaling is removed.
- T3-Response timeout—The maximum wait time for a response before removing the connection.
- Tunnel—The maximum period of inactivity for the GTP tunnel before it is torn down.

Step 7 Click the **IMSI Prefix Filtering** tab and configure IMSI prefix filtering if desired.

By default, GTP inspection does not check for valid Mobile Country Code (MCC)/Mobile Network Code (MNC) combinations. If you configure IMSI prefix filtering, the MCC and MNC in the IMSI of the received packet is compared with the configured MCC/MNC combinations. The system then takes one of the following actions based on the **Drop** option:

- **Drop** not selected (default)—The packet is dropped if it does not match any of the combinations.
- **Drop** selected—The packet is dropped if it does match at least one combination.

The Mobile Country Code is a non-zero, three-digit value; add zeros as a prefix for one- or two-digit values. The Mobile Network Code is a two- or three-digit value.

Add all MCC and MNC combinations you want to either permit or to drop. By default, the ASA 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*.

Step 8 Click the **Inspections** tab and define the specific inspections you want to implement based on traffic characteristics.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). Then, configure the criterion:

- **Access Point Name**—Matches the access point name against the specified regular expression or regular expression class. By default, all messages with valid access point names are inspected and any name is allowed.
- **Message ID**—Matches the message ID, from 1 to 255. You can specify one value or a range of values. You must specify whether the message is for GTPv1 (which includes GTPv0) or GTPv2. By default, all valid message IDs are allowed.
- **Message Length**—Matches messages where the length of the UDP payload is between the specified minimum and maximum length.
- **Version**—Matches the GTP version, from 0 to 255. You can specify one value or a range of values. By default all GTP versions are allowed.

- **MSISDN**—Matches the Mobile Station International Subscriber Directory Number (MSISDN) information element in the Create PDP Context request, Create session request, and Modify Bearer Response messages against the specified regular expression or regular expression class. The regular expression can identify a specific MSISDN, or a range of MSISDNs based on the first x number of digits. MSISDN filtering is supported for GTPv1 and GTPv2 only.
- **Selection Mode**—Matches the Selection Mode information element in the Create PDP Context request. The selection mode specifies the origin of the Access Point Name (APN) in the message, and can be one of the following. Selection Mode filtering is supported for GTPv1 and GTPv2 only.
  - 0—Verified. The APN was provided by the mobile station or network, and the subscription is verified.
  - 1—Mobile Station. The APN was provided by the mobile station, and the subscription is not verified.
  - 2—Network. The APN was provided by the network, and the subscription is not verified.
  - 3—Reserved, not used.

c) For Message ID matching, choose whether to drop the packet or to apply a rate limit in packets per second. The action for all other matches is to drop the packet. For all matches, you can choose whether to enable logging.

d) Click **OK** to add the inspection. Repeat the process as needed.

Step 9 Click the **Anti-Replay Protection** tab and configure the anti-replay options.

- **Enable Data Packet Replay Window**—Whether to enable anti-replay by specifying a sliding window for GTP-U messages. The size of the sliding window is in number of messages and can be 128, 256, 512, or 1024. As valid messages appear, the window moves to the new sequence numbers. Sequence numbers are in the range 0-65535, wrapping when they reach the maximum, and they are unique per PDP context. Messages are considered valid if their sequence numbers are within the window. Anti-replay helps prevent session hijacking or DoS attacks, which can occur when a hacker captures GTP data packets and replays them.

Step 10 Click the **User-Spoofing** tab and configure the anti-spoofing options.

- **GTP Header Check**—Whether to check that the inner payload of a GTP data packet is a valid IP packet, and drop the packet if it has a non-IP header. You must select this option to implement anti-spoofing.
- **Anti-Spoofing**—Whether to check that the mobile user IP address in the IP header of the inner payload matches the IP address assigned in GTP control messages such as Create Session Response, and drop the message if the IP addresses do not match. It is possible for hackers to pretend (spoof) that they are another customer by using another IP address than the one assigned through GTP-C. Anti-spoofing checks whether the GTP-U address used is actually the one which was assigned using GTP-C. This check supports IPv4, IPv6, and IPv4v6 PDN Types.

  If the mobile station gets its address using DHCP, the end-user IP address in GTPv2 is 0.0.0.0 (IPv4) or *prefix*::0 (IPv6), so in this case, the system updates the end-user IP address with the first IP address found in the inner packets. You can change the default behavior for DHCP-obtained addresses using the following keywords:

  - **GTPV2-DHCP-ByPass**—Do not update the 0.0.0.0 or *prefix*::0 address. Instead, allow packets where the end-user IP address is 0.0.0.0 or *prefix*::0. This option bypasses the anti-spoofing check when DHCP is used to obtain the IP address.

- **GTPV2-DHCP-DROP**—Do not update the 0.0.0.0 or *prefix*::0 address. Instead, drop all packets where the end-user IP address is 0.0.0.0 or *prefix*::0. This option prevents access for users that use DHCP to obtain the IP address.

Step 11 Click the **Location-Logging** tab and configure the location logging options.

- **Location Logging**—Whether to log the location of subscribers to track location changes for mobile stations. Tracking location changes can help you identify fraudulent roaming charges. When you enable location logging, the system generates syslog messages for new (message 324010) or changed (message 324011) location for each International Mobile Subscriber Identity (IMSI).

  Select the **Cell-ID** option if you want the log messages to include the Cell Global Identification (CGI) or E-UTRAN Cell Global Identifier (ECGI) where the user currently is registered.

Step 12 Click **OK** in the GTP Inspect Map dialog box.

You can now use the inspection map in a GTP inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure the Mobile Network Inspection Service Policy , on page 399.

Configure an SCTP Inspection Policy Map

To apply alternative actions to SCTP traffic based on the application-specific payload protocol identifier (PPID), such as rate limiting, create an SCTP inspection policy map to be used by the service policy.

Note PPIDs are in data chunks, and a given packet can have multiple data chunks or even a control chunk. If a packet includes a control chunk or multiple data chunks, the packet will not be dropped even if the assigned action is drop. For example, if you configure an SCTP inspection policy map to drop PPID 26, and a PPID 26 data chunk is combined in a packet with a Diameter PPID data chunk, that packet will not be dropped.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **SCTP**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map and click **Edit**.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 Drop, rate limit, or log traffic based on the PPID in SCTP data chunks.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose the match type for the criteria: **Match** (traffic must match the PPID) or **No Match** (traffic must not match the PPID).

For example, if you select No Match is on the Diameter PPID, then all PPIDs except Diameter are excluded from the class map.

c) Choose the **Minimum Payload PID** and optionally, the **Maximum Payload PID** to match.

You can enter PPIDs by name or number (0-4294967295). Click the **...** button in each field to select from a list of PPIDs. If you select a maximum PPID, then the match applies to the range of PPIDs

You can find the current list of SCTP PPIDs at http://www.iana.org/assignments/sctp-parameters/sctp-parameters.xhtml#sctp-parameters-25.

d) Choose whether to drop (and log), log, or rate limit (in kilobits per second, kbps) the matching packets.

e) Click **OK** to add the inspection. Repeat the process as needed.

Step 5 Click **OK** in the SCTP Inspect Map dialog box.

You can now use the inspection map in an SCTP inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure the Mobile Network Inspection Service Policy , on page 399.

Configure a Diameter Inspection Policy Map

You can create a Diameter inspection policy map to filter on various Diameter protocol elements. You can then selectively drop or log connections.

To configure Diameter message filtering, you must have a good knowledge of these protocol elements as they are defined in RFCs and technical specifications. For example, the IETF has a list of registered applications, command codes, and attribute-value pairs at http://www.iana.org/assignments/aaa-parameters/aaa-parameters.xhtml, although Diameter inspection does not support all listed items. See the 3GPP web site for their technical specifications.

You can optionally create a Diameter inspection class map to define the message filtering criteria for Diameter inspection. The other option is to define the filtering criteria directly in the Diameter inspection policy map. The difference between creating a class map and defining the filtering criteria directly in the inspection map is that you can create more complex match criteria and you can reuse class maps. Although this procedure explains inspection maps, the matching criteria used in class maps are the same as those explained in the step relating to the **Inspection** tab. You can configure Diameter class maps by selecting **Configuration** > **Firewall** > **Objects** > **Class Maps** > **Diameter**, or by creating them while configuring the inspection map.

Tip You can configure inspection maps while creating service policies, in addition to the procedure explained below. The contents of the map are the same regardless of how you create it.

Before you begin

Some traffic matching options use regular expressions for matching purposes. If you intend to use one of those techniques, first create the regular expression or regular expression class map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **Diameter**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map and click **Edit**. to view its contents.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 Click the **Parameters** tab and choose the desired options. whether you want to log messages that include unsupported Diameter elements.

- **Unsupported Parameters**—Whether you want to log messages that include unsupported Diameter elements. You can log unsupported **Application ID**, **Command Code**, or **Attribute Value Pair** elements.
- **Strict Diameter Validation Parameters**—Enables strict Diameter protocol conformance to RFC 6733. By default, inspection ensures that Diameter frames comply with the RFC. You can add session-related message validation and state machine validation.

Step 5 Click the **Inspections** tab and define the specific inspections you want to implement based on traffic characteristics.

You can define traffic matching criteria based on Diameter class maps, by configuring matches directly in the inspection map, or both.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose **Single Match** to define the criterion directly, or **Multiple Match**, in which case you select the Diameter class map that defines the criteria.

c) If you are defining the criterion here, choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). Then, configure the criterion as follows:

- **Application ID**—Enter the Diameter application name or number (0-4294967295). If there is a range of consecutively-numbered applications that you want to match, you can include a second ID. You can define the range by application name or number, and it applies to all the numbers between the first and second IDs.

These applications are registered with the IANA. Following are the core supported applications, but you can filter on other applications.

- **3gpp-rx-ts29214** (16777236)
- **3gpp-s6a** (16777251)
- **3gpp-s9** (16777267)
- **common-message** (0). This is the base Diameter protocol.

- **Command Code**—Enter the Diameter command code name or number (0-4294967295). If there is a range of consecutively-numbered command codes that you want to match, you can include a second code. You can define the range by command code name or number, and it applies to all the numbers between the first and second codes.

  For example, to match the Capability Exchange Request/Answer command code, CER/CEA, enter **cer-cea**.

- **Attribute Value Pair**—You can match the AVP by attribute only, a range of AVPs, or an AVP based on the value of the attribute. For the **AVP Begin Value**, you can specify the name of a custom AVP or one that is registered in RFCs or 3GPP technical specifications and is directly supported in the software. Click the **...** button in the field to pick from a list.

  If you want to match a range of AVP, specify the **AVP End Value** by number only. If you want to match an AVP by its value, you cannot specify a second code.

  You can further refine the match by specifying the optional **Vendor ID**, from 0-4294967295. For example, the 3GPP vendor ID is 10415, the IETF is 0.

  You can configure value-matching only if the data type of the AVP is supported. For example, you can specify an IP address for AVP that have the address data type. The list of AVP shows the data type for each. How you specify the value differs based on the AVP data type:

  - Diameter Identity, Diameter URI, Octet String—Select the regular expression or regular expression class objects to match these data types.
  - Address—Specify the IPv4 or IPv6 address to match. For example, 10.100.10.10 or 2001:DB8::0DB8:800:200C:417A.
  - Time—Specify the start and end dates and time. Both are required. Time is in 24-hour format.
  - Numeric—Specify a range of numbers. The valid number range depends on the data type:
    - Integer32: -2147483647 to 2147483647
    - Integer64: -9223372036854775807 to 9223372036854775807
    - Unsigned32: 0 to 4294967295
    - Unsigned64: 0 to 18446744073709551615
    - Float32: decimal point representation with 8 digit precision
    - Float64: decimal point representation with 16 digit precision

d) Choose the action to take for matching traffic: drop packet, drop connection, or log.
e) Click **OK** to add the inspection. Repeat the process as needed.

Step 6 Click **OK** in the Diameter Inspect Map dialog box.

You can now use the inspection map in a Diameter inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure the Mobile Network Inspection Service Policy , on page 399.

Create a Custom Diameter Attribute-Value Pair (AVP)

As new attribute-value pairs (AVP) are defined and registered, you can create custom Diameter AVP to define them and use them in your Diameter inspection policy map. You would get the information you need to create the AVP from the RFC or other source that defines the AVP.

Create custom AVP only if you want to use them in a Diameter inspection policy map or class map for AVP matching.

Procedure

Step 1 Select **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **Diameter AVP**.

Step 2 Click **Add** to create a new AVP.

When you edit an AVP, you can change the description only.

Step 3 Configure the following options:

- **Name**—The name of the custom AVP you are creating, up to 32 characters. You would refer to this name in a Diameter inspection policy map or class map when defining an attribute-value pair match.
- **Custom Code**—The custom AVP code value, from 256-4294967295. You cannot enter a code and vendor ID combination that is already defined in the system.
- **Data Type**—The data type of the AVP. You can define AVP of the following types. If the new AVP is of a different type, you cannot create a custom AVP for it.
  - Address (for IP addresses)
  - Diameter identity
  - Diameter uniform resource identifier (URI)
  - 32-bit floating point number
  - 64-bit floating point number
  - 32-bit integer
  - 64-bit integer
  - Octet string
  - Time

- 32-bit unsigned integer
- 64-bit unsigned integer

- **Vendor ID**—(Optional.) The ID number of the vendor who defined the AVP, from 0-4294967295. For example, the 3GPP vendor ID is 10415, the IETF is 0.
- **Description**—(Optional.) A description of the AVP, up to 80 characters.

Step 4 Click **OK**.

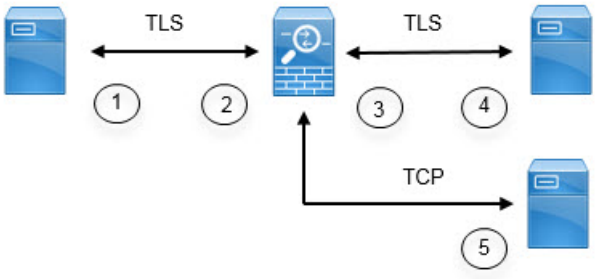
Inspecting Encrypted Diameter Sessions

If a Diameter application uses encrypted data over TCP, inspection cannot see inside the packets to implement your message filtering rules. Thus, if you create filtering rules, and you want them to also apply to encrypted TCP traffic, you must configure a TLS proxy. You also need a proxy if you want strict protocol enforcement on encrypted traffic. This configuration does not apply to SCTP/DTLS traffic.

The TLS proxy acts as a man-in-the-middle. It decrypts traffic, inspects it, then encrypts it again and sends it to the intended destination. Thus, both sides of the connection, the Diameter server and Diameter client, must trust the ASA, and all parties must have the required certificates. You must have a good understanding of digital certificates to implement TLS proxy. Please read the chapter on digital certificates in the ASA general configuration guide.

The following illustration shows the relationship among the Diameter client and server, and the ASA, and the certification requirements to establish trust. In this model, a Diameter client is an MME (Mobility Management Entity), not an end user. The CA certificate on each side of a link is the one used to sign the certificate on the other side of the link. For example, the ASA proxy TLS server CA certificate is the one used to sign the Diameter/TLS client certificate.

Figure 49: Diameter TLS Inspection



1	Diameter TLS client (MME) • Client identity certificate • CA certificate used to sign the ASA TLS proxy server's identity certificate	2	ASA proxy TLS server • Server identity certificate • CA certificate used to sign the Diameter TLS client's identity certificate

3	ASA proxy TLS client • Client identity (static or LDC) certificate • CA certificate used to sign the Diameter TLS server identity certificate	4	Diameter TLS server (full proxy) • Server identity certificate • CA certificate used to sign the ASA proxy TLS client's identity certificate
5	Diameter TCP server (TLS offload).	—	—

You have the following options for configuring TLS proxy for Diameter inspection:

- Full TLS proxy—Encrypt traffic between the ASA and Diameter clients and the ASA and Diameter server. You have the following options for establishing the trust relationship with the TLS server:
  - Use a static proxy client trustpoint. The ASA presents the same certificate for every Diameter client when communicating with the Diameter server. Because all clients look the same, the Diameter server cannot provide differential services per client. On the other hand, this option is faster than the LDC method.
  - Use local dynamic certificates (LDC). With this option, the ASA presents unique certificates per Diameter client when communicating with the Diameter server. The LDC retains all fields from the received client identity certificate except its public key and a new signature from the ASA. This method gives the Diameter server better visibility into client traffic, which makes it possible to provide differential services based on client certificate characteristics.
- TLS offload—Encrypt traffic between the ASA and Diameter client, but use a clear-text connection between the ASA and Diameter server. This option is viable if the Diameter server is in the same data center as the ASA, where you are certain that the traffic between the devices will not leave the protected area. Using TLS offload can improve performance, because it reduces the amount of encryption processing required. It should be the fastest of the options. The Diameter server can apply differential services based on client IP address only.

All three options use the same configuration for the trust relationship between the ASA and Diameter clients.

Note TLS proxy uses TLSv1.0 - 1.2. You can configure the TLS version and the cipher suite.

The following topics explain how to configure TLS proxy for Diameter inspection.

Configure Server Trust Relationship with Diameter Clients

The ASA acts as a TLS proxy server in relation to the Diameter clients. To establish the mutual trust relationship:

- You need to import the Certificate Authority (CA) certificate used to sign the ASA's server certificate into the Diameter client. This might be in the client's CA certificate store or some other location that the client uses. Consult the client documentation for exact details on certificate usage.
- You need to import the CA certificate used to sign the Diameter TLS client's certificate so the ASA can trust the client.

The following procedure explains how to import the CA certificate used to sign the Diameter client's certificate, and import an identity certificate to use for the ASA TLS proxy server. Instead of importing an identity certificate, you could create a self-signed certificate on the ASA. Alternatively, you can import these certificates when you create the TLS proxy.

Procedure

Step 1 Import the CA certificate that is used to sign the Diameter client's certificate into an ASA trustpoint.

This step allows the ASA to trust the Diameter clients.

a) Select **Configuration** > **Firewall** > **Advanced** > **Certificate Management** > **CA Certificates**.
b) Click **Add** and enter a name for the trustpoint. For example, **diameter-clients**.
c) Add the certificate.

You can import the certificate from a file, paste it in PEM format, or use SCEP to import it.

d) Click **Install Certificate**.

Step 2 Import the certificate and create a trustpoint for the ASA proxy server's identity certificate and keypair.

This step allows the Diameter clients to trust the ASA.

a) Select **Configuration** > **Firewall** > **Advanced** > **Certificate Management** > **Identity Certificates**.
b) Click **Add** and enter a name for the trustpoint. For example, **tls-proxy-server-tp**.
c) Select **Import the identity certificate from a file**, enter the decryption passphrase, and select the file (in pkcs12 format).

Alternatively, you can create a new certificate.

d) Click **Add Certificate**.

Configure Full TLS Proxy with Static Client Certificate for Diameter Inspection

If the Diameter server can accept the same certificate for all clients, you can set up a static client certificate for the ASA to use when communicating with the Diameter server.

With this configuration, you need to establish the mutual trust relationship between the ASA and clients (as explained in Configure Server Trust Relationship with Diameter Clients, on page 391), and the ASA and Diameter server. Following are the ASA and Diameter server trust requirements.

- You need to import the CA certificate used to sign the Diameter Server's identity certificate so the ASA can validate the server's identity certificate during the TLS handshake.
- You need to import the client certificate, one that the Diameter server also trusts. If the Diameter server does not already trust the certificate, import the CA certificate used to sign it into the server. Consult the Diameter server's documentation for details.

Procedure

Step 1 Select **Configuration** > **Firewall** > **Unified Communications** > **TLS Proxy**.

Step 2 Click **Add**.

Step 3 Give the TLS proxy a name, for example, **diameter-tls-static-proxy**. Click **Next**.

Step 4 Select the TLS server proxy identity certificate that you added in Configure Server Trust Relationship with Diameter Clients, on page 391. Click **Next**.

If you have not already created the identity certificate, you can click **Manage** to add it. You can also install the Diameter client's CA certificate by clicking **Install TLS Server's Certificate**.

Optionally, you can define the security algorithms (ciphers) that the server can use by moving them from the available algorithms to the active algorithms list. If you do not specify ciphers, the default system ciphers are used.

Note
For testing purposes, or if you are certain that you can trust the Diameter clients, you can skip this step and deselect **Enable client authentication during TLS Proxy handshake** in the TLS proxy configuration.

Step 5 Select **Specify the proxy certificate for TLS client** and do the following:

a) Select the certificate for the ASA TLS proxy client.

If you have not already added the certificate, click **Manage** and add it now.

b) If you have not already added the CA certificate that was used to sign the Diameter server's certificate, click **Install TLS Client's Certificate** and add it.

c) (Optional.) Define the security algorithms (ciphers) that the client can use by moving them from the available algorithms to the active algorithms list.

If you do not define the ciphers the TLS proxy can use, the proxy uses the global cipher suite defined by the **Configuration** > **Device Management** > **Advanced** > **SSL Settings** encryption settings. By default, the global cipher level is medium, which means all ciphers are available except for NULL-SHA, DES-CBC-SHA, and RC4-MD5. Specify the TLS proxy-specific ciphers only if you want to use a different suite than the one generally available on the ASA.

d) Click **Next**.

Step 6 Click **Finish**, then click **Apply**.

What to do next

You can now use the TLS proxy in Diameter inspection. See Configure the Mobile Network Inspection Service Policy , on page 399.

Configure Full TLS Proxy with Local Dynamic Certificates for Diameter Inspection

If the Diameter server needs unique certificates for each client, you can configure the ASA to generate local dynamic certificates (LDC). These certificates exist for the duration of the client's connection and are then destroyed.

With this configuration, you need to establish the mutual trust relationship between the ASA and clients (as explained in Configure Server Trust Relationship with Diameter Clients, on page 391), and the ASA and Diameter server. The configuration is similar to the one described in Configure Full TLS Proxy with Static Client Certificate for Diameter Inspection, on page 392, except instead of importing a Diameter client certificate, you set up the LDC on the ASA. Following are the ASA and Diameter server trust requirements.

- You need to import the CA certificate used to sign the Diameter Server's identity certificate so the ASA can validate the server's identity certificate during the TLS handshake.
- You need to create the LDC trustpoint. You need to export the LDC server's CA certificate and import it into the Diameter server. The export step is explained below. Consult the Diameter server's documentation for information on importing certificates.

Procedure

Step 1 Select **Configuration** > **Firewall** > **Unified Communications** > **TLS Proxy**.

Step 2 Click **Add**.

Step 3 Give the TLS proxy a name, for example, **diameter-tls-ldc-proxy**.

Step 4 Select the TLS server proxy identity certificate that you added in Configure Server Trust Relationship with Diameter Clients, on page 391. Click **Next**.

If you have not already created the identity certificate, you can click **Manage** to add it. You can also install the Diameter client's CA certificate by clicking **Install TLS Server's Certificate**.

Optionally, you can define the security algorithms (ciphers) that the server can use by moving them from the available algorithms to the active algorithms list. If you do not specify ciphers, the default system ciphers are used.

Note
For testing purposes, or if you are certain that you can trust the Diameter clients, you can skip this step and deselect **Enable client authentication during TLS Proxy handshake** in the TLS proxy configuration.

Step 5 Select **Specify the internal Certificate Authority to sign for local dynamic certificates** and do the following (ignore any text related to IP phones).

This procedure assumes you are creating a new certificate and key. If you have already created the needed certificate and key, select them and move to the security algorithms step.

a) For Local Dynamic Certificate Key Pair, click **New**. (You might need to resize the dialog box to see the button.)
b) Create a general purpose RSA certificate with a new key pair name, such as **ldc-signer-key**. Click **Generate Now** to create the key.

   You are returned to the Manage Identity Certificates dialog box.

c) Select **Certificate** and click **Manage** to create the certificate and key for the ASA TLS proxy client.
d) Click **Add** in the Manage Identity Certificates dialog box.
e) Give the trustpoint a name, such as **ldc-server**.
f) Select **Add a new identity certificate**.
g) For **Key Pair**, select the same key you created for the local dynamic certificate key.
h) For **Certificate Subnet DN**, select the Distinguished Name attributes that you need.

   The device's common name is the default. Check whether the Diameter application has specific requirements for the subject name.

i) Select **Generate self-signed certificate**. This is required.
j) Select **Act as a local certificate authority and issue dynamic certificates to TLS Proxy**. This option make this certificate an LDC issuer.

k) Click **Add Certificate**.

You are returned to the Manage Identity Certificates dialog box.

l) Select the certificate you just created and click **Export**.

You need to export the certificate so that you can import it into the Diameter server. Specify a file name and PEM format, and click **Export Certificate**.

You are returned to the Manage Identity Certificates dialog box.

m) With the certificate still selected, click **OK**.

You are returned to the TLS Proxy wizard. If the certificate is not selected in the Certificate field, select it now.

n) (Optional.) Define the security algorithms (ciphers) that the client can use by moving them from the available algorithms to the active algorithms list.

If you do not define the ciphers the TLS proxy can use, the proxy uses the global cipher suite defined by the **Configuration** > **Device Management** > **Advanced** > **SSL Settings** encryption settings. By default, the global cipher level is medium, which means all ciphers are available except for NULL-SHA, DES-CBC-SHA, and RC4-MD5. Specify the TLS proxy-specific ciphers only if you want to use a different suite than the one generally available on the ASA.

o) Click **Next**.

Step 6 Click **Finish**, then click **Apply**.

Step 7 You can now import the LDC CA certificate into the Diameter server. Consult the Diameter server's documentation for the procedure. Note that the data is in Base64 format. If your server requires binary or DER format, you will need to use OpenSSL tools to convert formats.

What to do next

You can now use the TLS proxy in Diameter inspection. See Configure the Mobile Network Inspection Service Policy , on page 399.

Configure TLS Proxy with TLS Offload for Diameter Inspection

If you are certain the network path between the ASA and Diameter server is secure, you can avoid the performance cost of encrypting data between the ASA and server. With TLS offload, the TLS proxy encrypts/decrypts sessions between the Diameter client and the ASA, but uses clear text with the Diameter server.

With this configuration, you need to establish the mutual trust relationship between the ASA and clients only, which simplifies the configuration. Before doing the following procedure, complete the steps in Configure Server Trust Relationship with Diameter Clients, on page 391.

Procedure

Step 1 Select **Configuration** > **Firewall** > **Unified Communications** > **TLS Proxy**.

Step 2 Click **Add**.

Step 3 Give the TLS proxy a name, for example, **diameter-tls-offload-proxy**.

Step 4 Select the TLS server proxy identity certificate that you added in Configure Server Trust Relationship with Diameter Clients, on page 391. Click **Next**.

If you have not already created the identity certificate, you can click **Manage** to add it. You can also install the Diameter client's CA certificate by clicking **Install TLS Server's Certificate**.

Optionally, you can define the security algorithms (ciphers) that the server can use by moving them from the available algorithms to the active algorithms list. If you do not specify ciphers, the default system ciphers are used.

Note
For testing purposes, or if you are certain that you can trust the Diameter clients, you can skip this step and deselect **Enable client authentication during TLS Proxy handshake** in the TLS proxy configuration.

Step 5 Select **Configure the proxy client to use clear text to communicate with the remote TCP client**, and click **Next**.

Step 6 Click **Finish**, then click **Apply**.

Step 7 Because the Diameter ports differ for TCP and TLS, configure a NAT rule to translate the TCP port to the TLS port for traffic going from the Diameter server to the client.

Create an object NAT rule for each Diameter server.

a) Select **Configuration** > **Firewall** > **NAT**.
b) Click **Add** > **Object NAT Rule**.
c) Configure the basic properties:

- **Name**—The object name, for example, DiameterServerA.
- **Type** (for the object)—Select **Host**.
- **IP Version**—IPv4 or IPv6 as appropriate.
- **IP Address**—The IP address of the Diameter server, for example, 10.100.10.10.
- **Add Automatic Address Translation**—Ensure you select this option.
- **Type** (for the NAT rule)—Select **Static**.
- **Translated Addr**—The IP address of the Diameter server. This would be the same as the IP Address for the object, for example, 10.100.10.10.

d) Click **Advanced** and configure the following **Interface** and **Service** options:

- **Source Interface**—Select the interface that connects to the Diameter server.
- **Destination Interface**—Select the interface that connects to the Diameter client.
- **Protocol**—Select **TCP**.
- **Real Port**—Enter 3868, which is the default Diameter TCP port number.
- **Mapped Port**—Enter 5868, which is the default Diameter TLS port number.

e) Click **OK**, then click **OK** again in the Add Network Object dialog box.

What to do next

You can now use the TLS proxy in Diameter inspection. See Configure the Mobile Network Inspection Service Policy , on page 399.

Configure an M3UA Inspection Policy Map

Use an M3UA inspection policy map to configure access control based on point codes. You can also drop and rate limit messages by class and type.

The default point code format is ITU. If you use a different format, specify the required format in the policy map.

If you do not want to apply policy based on point code or message class, you do not need to configure an M3UA policy map. You can enable inspection without a map.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **M3UA**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map and click **Edit** to edit the map.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 Click the **Parameters** tab and configure the desired options:

- **SS7**—The variant of SS7 used in your network: ITU, ANSI, Japan, China. This option determines the valid format for point codes. After you configure the option and deploy an M3UA policy, you cannot change it unless you first remove the policy. The default variant is ITU.
- **Enable M3UA Application Server Process (ASP) state validation**—Whether to perform strict application server process (ASP) state validation. The system maintains the ASP states of M3UA sessions and allows or drops ASP messages based on the validation result. If you do not enable strict ASP state validation, all ASP messages are forwarded uninspected. Strict ASP state checking is required if you want stateful failover or if you want to operate within a cluster. However, strict ASP state checking works in Override mode only, it does not work if you are running in Loadsharing or Broadcast mode (per RFC 4666). The inspection assumes there is one and only one ASP per endpoint.
- **Enforce Timeout** > **Endpoint**—The idle timeout to remove statistics for an M3UA endpoint, in hh:mm:ss format. To have no timeout, specify 0. The default is 30 minutes (00:30:00).
- **Enforce Timeout** > **Session**—The idle timeout to remove an M3UA session if you enable strict ASP state validation, in hh:mm:ss format. To have no timeout, specify 0. The default is 30 minutes (00:30:00). Disabling this timeout can prevent the system from removing stale sessions.
- **Message Tag Validation**—Whether to check and validate the content of certain fields for the specified message type. Messages that fail validation are dropped. Validation differs by message type. Select the messages you want to validate.

- Destination User Part Unavailable (DUPU)—The User/Cause field must be present, and it must contain only valid cause and user codes.
- Error—All mandatory fields must be present and contain only allowed values. Each error message must contain the required fields for that error code.
- Notify—The status type and status information fields must contain allowed values only.

Step 5 Click the **Inspections** tab and define the specific inspections you want to implement based on traffic characteristics.

a) Do any of the following:

- Click **Add** to add a new criterion.
- Select an existing criterion and click **Edit**.

b) Choose the match type for the criteria: **Match** (traffic must match the criterion) or **No Match** (traffic must not match the criterion). Then, configure the criterion:

- **Class ID**—Matches the M3UA message class and type. The following table lists the possible values. Consult M3UA RFCs and documentation for detailed information about these messages.

M3UA Message Class	Message ID Type
0 (Management Messages)	0-1
1 (Transfer Messages)	1
2 (SS7 Signaling Network Management Messages)	1-6
3 (ASP State Maintenance Messages)	1-6
4 (ASP Traffic Maintenance Messages)	1-4
9 (Routing Key Management Messages)	1-4

- **OPC**—Matches the originating point code, that is, the traffic source. Point code is in *zone-region-sp* format, where the possible values for each element depend on the SS7 variant:
  - **ITU**—Point codes are 14 bit in 3-8-3 format. The value ranges are [0-7]-[0-255]-[0-7].
  - **ANSI**—Point codes are 24 bit in 8-8-8 format. The value ranges are [0-255]-[0-255]-[0-255].
  - **Japan**—Point codes are 16 bit in 5-4-7 format. The value ranges are [0-31]-[0-15]-[0-127].
  - **China**—Point codes are 24 bit in 8-8-8 format. The value ranges are [0-255]-[0-255]-[0-255].
- **DPC**—Matches the destination point code. Point code is in *zone-region-sp* format, as explained for **OPC**.
- **Service Indicator**—Matches the service indicator number, 0-15. Following are the available service indicators. Consult M3UA RFCs and documentation for detailed information about these service indicators.
  - 0—Signaling Network Management Messages

- 1—Signaling Network Testing and Maintenance Messages
- 2—Signaling Network Testing and Maintenance Special Messages
- 3—SCCP
- 4—Telephone User Part
- 5—ISDN User Part
- 6—Data User Part (call and circuit-related messages)
- 7—Data User Part (facility registration and cancellation messages)
- 8—Reserved for MTP Testing User Part
- 9—Broadband ISDN User Part
- 10—Satellite ISDN User Part
- 11—Reserved
- 12—AAL type 2 Signaling
- 13—Bearer Independent Call Control
- 14—Gateway Control Protocol
- 15—Reserved

c) For Class ID matching, choose whether to drop the packet or to apply a rate limit in packets per second. The action for all other matches is to drop the packet. For all matches, you can choose whether to enable logging.
d) Click **OK** to add the inspection. Repeat the process as needed.

Step 6 Click **OK** in the M3UA Inspect Map dialog box.

You can now use the inspection map in an M3UA inspection service policy.

What to do next

You can now configure an inspection policy to use the map. See Configure the Mobile Network Inspection Service Policy , on page 399.

Configure the Mobile Network Inspection Service Policy

Inspections for the protocols used in mobile networks are not enabled in the default inspection policy, so you must enable them if you need these inspections. You can simply edit the default global inspection policy to add these inspections. You can alternatively create a new service policy as desired, for example, an interface-specific policy.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy**, and open a rule.

- To edit the default global policy, select the "inspection_default" rule in the Global folder and click **Edit**.
- To create a new rule, click **Add** > **Add Service Policy Rule**. Proceed through the wizard to the Rules page.
- If you have a mobile network inspection rule, or a rule to which you are adding these inspections, select it and click **Edit**.

Step 2 On the Rule Actions wizard page or tab, select the **Protocol Inspection** tab.

Step 3 (To change an in-use policy.) If you are editing any in-use policy to use a different inspection policy map, you must disable the inspections, and then re-enable them with the new inspection policy map name:

a) Uncheck the relevant already-selected check boxes: **GTP**, **SCTP**, **Diameter**, **M3UA**.
b) Click **OK**.
c) Click **Apply**.
d) Repeat these steps to return to the Protocol Inspections tab.

Step 4 Select the desired mobile network protocols: **GTP**, **SCTP**, **Diameter**, **M3UA**.

Step 5 If you want non-default inspection for one or more of these protocols, click **Configure** next to the options, and do the following:

a) Choose whether to use the default map or to use an inspection policy map that you configured. You can create the map at this time.
b) (Diameter only.) To enable Diameter inspection of encrypted messages, select **Enable Encrypted Traffic Inspection**, and select a TLS proxy to use for decryption.

Note
If you specify a TLS proxy for Diameter inspection, and you apply NAT port redirection to Diameter server traffic (for example, redirect server traffic from port 5868 to 3868), configure inspection globally or on the ingress interface only. If you apply the inspection to the egress interface, NATed Diameter traffic bypasses inspection.

c) Click **OK** in the Select Inspect Map dialog box.

Step 6 Click **OK** or **Finish** to save the service policy rule.

Configure RADIUS Accounting Inspection

RADIUS accounting inspection is not enabled by default. You must configure it if you want RADIUS accounting inspection.

Procedure

Step 1 Configure a RADIUS Accounting Inspection Policy Map, on page 401.

Step 2 Configure the RADIUS Accounting Inspection Service Policy, on page 402.

Configure a RADIUS Accounting Inspection Policy Map

You must create a RADIUS accounting inspection policy map to configure the attributes needed for the inspection.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **RADIUS Accounting**.

Step 2 Do one of the following:

- Click **Add** to add a new map.
- Select a map and click **Edit**.

Step 3 For new maps, enter a name (up to 40 characters) and description. When editing a map, you can change the description only.

Step 4 Click the **Host Parameters** tab and add the IP addresses of each RADIUS server or GGSN.

You can optionally include a secret key so that the ASA can validate the message. Without the key, only the IP address is checked. The ASA receives a copy of the RADIUS accounting messages from these hosts.

Step 5 Click the **Other Parameters** tab and configure the desired options.

- **Send responses to the originator of the RADIUS accounting message**—xxWhether to mask the banner from the ESMTP server.
- **Enforce user timeout**—Whether to implement an idle timeout for users, and the timeout value. The default is one hour.
- **Enable detection of GPRS accounting**—Whether to implement GPRS over-billing protection. The ASA checks for the 3GPP VSA 26-10415 attribute in the Accounting-Request Stop and Disconnect messages in order to properly handle secondary PDP contexts. If this attribute is present, then the ASA tears down all connections that have a source IP matching the User IP address on the configured interface.
- **Validate Attribute**—Additional criteria to use when building a table of user accounts when receiving Accounting-Request Start messages. These attributes help when the ASA decides whether to tear down connections.

  If you do not specify additional attributes to validate, the decision is based solely on the IP address in the Framed IP Address attribute. If you configure additional attributes, and the ASA receives a start accounting message that includes an address that is currently being tracked, but the other attributes to validate are different, then all connections started using the old attributes are torn down, on the assumption that the IP address has been reassigned to a new user.

  Values range from 1-191, and you can enter the command multiple times. For a list of attribute numbers and their descriptions, see http://www.iana.org/assignments/radius-types.

Step 6 Click **OK**.

You can now use the inspection map in a RADIUS accounting inspection service policy.

Configure the RADIUS Accounting Inspection Service Policy

RADIUS accounting inspection is not enabled in the default inspection policy, so you must enable it if you need this inspection. Because RADIUS accounting inspection is for traffic directed to the ASA, you must configure it as a management inspection rule rather than a standard rule.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy**, and open a rule.

- To create a new rule, click **Add > Add Management Service Policy Rule**. Proceed through the wizard to the Rules page.
- If you have a RADIUS accounting inspection rule, or a management rule to which you are adding RADIUS accounting inspection, select it, click **Edit**, and click the **Rule Actions** tab.

Step 2 (To change an in-use policy) If you are editing any in-use policy to use a different inspection policy map, you must disable the RADIUS accounting inspection, and then re-enable it with the new inspection policy map name:

a) Select **None** for the RADIUS Accounting map.
b) Click **OK**.
c) Click **Apply**.
d) Repeat these steps to return to the Protocol Inspections tab.

Step 3 Choose the desired **RADIUS Accounting Map**. You can create the map at this time. For detailed information, see Configure a RADIUS Accounting Inspection Policy Map, on page 401.

Step 4 Click **OK** or **Finish** to save the management service policy rule.

Monitoring Mobile Network Inspection

The following topics explain how to monitor mobile network inspection.

Monitoring GTP Inspection

To display the GTP configuration, enter the **show service-policy inspect gtp** command in privileged EXEC mode. Select **Tools** > **Command Line Interface** to enter commands.

Use the **show service-policy inspect gtp statistics** command to show the statistics for GTP inspection. The following is sample output:

```
firewall(config)# 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                  67     total_dropped             1
  signalling_msg_dropped            1     data_msg_dropped          0
  signalling_msg_forwarded         67     data_msg_forwarded        0
  total created_pdp                33     total deleted_pdp         32
  total created_pdpmcb             31     total deleted_pdpmcb      30
  total dup_sig_mcbinfo             0     total dup_data_mcbinfo    0
  no_new_sgw_sig_mcbinfo            0     no_new_sgw_data_mcbinfo   0
  pdp_non_existent                  1
```

You can get statistics for a specific GTP endpoint by entering the IP address on the **show service-policy inspect gtp statistics** *ip_address* command.

```
firewall(config)# show service-policy inspect gtp statistics 10.9.9.9
1 in use, 1 most used, timeout 0:30:00
 GTP GSN Statistics for 10.9.9.9, Idle 0:00:34, restart counter 0
  Tunnels Active                     0
  Tunnels Created                    1
  Tunnels Destroyed                  0
  Total Messages Received            1
                         Signalling Messages        Data Messages
  total received                    1                          0
  dropped                           0                          0
  forwarded                         1                          0
```

Use the **show service-policy inspect gtp pdp-context** command to display PDP context-related information. For GTPv2, this is the bearer context. For example:

```
ciscoasa(config)# show service-policy inspect gtp pdp-context
4 in use, 5 most used

Version v1,   TID 050542012151705f,  MS Addr 2005:a00::250:56ff:fe96:eec,
SGSN Addr 10.0.203.22,      Idle 0:52:01,   Timeout 3:00:00,   APN ssenoauth146

Version v2,   TID 0505420121517056,  MS Addr 100.100.100.102,
SGW Addr 10.0.203.24,      Idle 0:00:05,   Timeout 3:00:00,   APN ssenoauth146

Version v2,   TID 0505420121517057,  MS Addr 100.100.100.103,
SGW Addr 10.0.203.25,      Idle 0:00:04,   Timeout 3:00:00,   APN ssenoauth146

Version v2,   TID 0505420121517055,  MS Addr 100.100.100.101,
SGW Addr 10.0.203.23,      Idle 0:00:06,   Timeout 3:00:00,   APN ssenoauth146

ciscoasa(config)# show service-policy inspect gtp pdp-context detail
1 in use, 1 most used

Version v1,   TID 050542012151705f,  MS Addr 2005:a00::250:56ff:fe96:eec,
SGSN Addr 10.0.203.22,      Idle 0:06:14,   Timeout 3:00:00,   APN ssenoauth146

    user_name (IMSI):  50502410121507    MS address: 2005:a00::250:56ff:fe96:eec
    nsapi: 5                linked nsapi: 5
    primary pdp: Y         sgsn is Remote
    sgsn_addr_signal: 10.0.203.22    sgsn_addr_data: 10.0.203.22
    ggsn_addr_signal: 10.0.202.22    ggsn_addr_data: 10.0.202.22
    sgsn control teid:     0x00000001    sgsn data teid:        0x000003e8
    ggsn control teid:     0x000f4240    ggsn data teid:        0x001e8480
    signal_sequence:              18     state:    Ready
```

...

The PDP or bearer context is identified by the tunnel ID (TID), which is a combination of the values for IMSI and NSAPI (GTPv0-1) or IMSI and EBI (GTPv2). A GTP tunnel is defined by two associated contexts in different GSN or SGW/PGW nodes and is identified with a Tunnel ID. A GTP tunnel is necessary to forward packets between an external packet data network and a mobile subscriber (MS) user.

Monitoring SCTP

You can use the following commands to monitor SCTP. Select **Tools** > **Command Line Interface** to enter these commands.

- **show service-policy inspect sctp**

  Displays SCTP inspection statistics. The sctp-drop-override counter increments each time a PPID is matched to a drop action, but the packet was not dropped because it contained data chunks with different PPIDs. For example:

```
ciscoasa# show service-policy inspect sctp
Global policy:
  Service-policy: global_policy
    Class-map: inspection_default
      Inspect: sctp sctp, packet 153302, lock fail 0, drop 20665, reset-drop 0,
5-min-pkt-rate 0 pkts/sec, v6-fail-close 0, sctp-drop-override 4910
        Match ppid 30 35
           rate-limit 1000 kbps, chunk 2354, dropped 10, bytes 21408, dropped-bytes 958

        Match: ppid 40
           drop, chunk 5849
        Match: ppid 55
           log, chunk 9546
```

- **show sctp** [**detail**]

  Displays current SCTP cookies and associations. Add the **detail** keyword to see detailed information about SCTP associations. The detailed view also shows information about multi-homing, multiple streams, and fragment reassembly.

```
ciscoasa# show sctp

AssocID: 71adeb15
Local:  192.168.107.12/50001 (ESTABLISHED)
Remote: 192.168.108.122/2905 (ESTABLISHED)
Secondary Conn List:
    192.168.108.12(192.168.108.12):2905 to 192.168.107.122(192.168.107.122):50001
    192.168.107.122(192.168.107.122):50001 to 192.168.108.12(192.168.108.12):2905
    192.168.108.122(192.168.108.122):2905 to 192.168.107.122(192.168.107.122):50001
    192.168.107.122(192.168.107.122):50001 to 192.168.108.122(192.168.108.122):2905
    192.168.108.12(192.168.108.12):2905 to 192.168.107.12(192.168.107.12):50001
    192.168.107.12(192.168.107.12):50001 to 192.168.108.12(192.168.108.12):2905
```

- **show conn protocol sctp**

  Displays information about current SCTP connections.

- **show local-host** [**connection sctp** *start*[*-end*]]

Displays information on hosts making SCTP connections through the ASA, per interface. Add the **connection sctp** keyword to see only those hosts with the specified number or range of SCTP connections.

- **show traffic**

  Displays SCTP connection and inspection statistics per interface if you enable the **sysopt traffic detailed-statistics** command.

Monitoring Diameter

You can use the following commands to monitor Diameter. Select **Tools** > **Command Line Interface** to enter these commands.

- **show service-policy inspect diameter**

  Displays Diameter inspection statistics. For example:

```
ciscoasa# show service-policy inspect diameter
Global policy:
  Service-policy: global_policy
    Class-map: inspection_default
      Inspect: Diameter Diameter_map, packet 0, lock fail 0, drop 0, -drop 0,
5-min-pkt-rate 0 pkts/sec, v6-fail-close 0
       Class-map: log_app
           Log: 5849
        Class-map: block_ip
            drop-connection: 2
```

- **show diameter**

  Displays state information for each Diameter connection. For example:

```
ciscoasa# show diameter
Total active diameter sessions: 5
Session 3638
       ==========
       ref_count: 1 val = .; 1096298391; 2461;
           Protocol : diameter Context id : 0
           From inside:211.1.1.10/45169 to outside:212.1.1.10/3868
...
```

- **show conn detail**

  Displays connection information. Diameter connections are marked with the Q flag.

- **show tls-proxy**

  Displays information about the TLS proxy if you use one in Diameter inspection.

Monitoring M3UA

You can use the following commands to monitor M3UA. Select **Tools** > **Command Line Interface** to enter these commands.

- **show service-policy inspect m3ua drops**

Displays drop statistics for M3UA inspection.

- **show service-policy inspect m3ua endpoint** [*IP_address*]

  Displays statistics for M3UA endpoints. You can specify an endpoint IP address to see information for a specific endpoint. For high availability or clustered systems, the statistics are per unit, they are not synchronized across units. For example:

```
ciscoasa# sh service-policy inspect m3ua endpoint
 M3UA Endpoint Statistics for 10.0.0.100, Idle : 0:00:06 :
                   Forwarded        Dropped          Total Received
 All Messages      21                5                 26
 DATA Messages     9                 5                 14
 M3UA Endpoint Statistics for 10.0.0.110, Idle : 0:00:06 :
                   Forwarded        Dropped          Total Received
 All Messages      21                8                 29
 DATA Messages     9                 8                 17
```

- **show service-policy inspect m3ua session**

  Displays information about M3UA sessions if you enable strict application server process (ASP) state validation. Information includes source association ID, whether the session is single or double exchange, and in clustering, whether it is a cluster owner session or a backup session. In a cluster with 3 or more units, you might see stale backup sessions if a unit leaves and then returns to the cluster. These stale sessions are removed when they time out, unless you disabled session timeout.

```
Ciscoasa# show service-policy inspect m3ua session
0 in use, 0 most used
Flags: o - cluster owner session, b - cluster backup session
       d - double exchange      , s - single exchange
AssocID: cfc59fbe in Down state, idle:0:00:05, timeout:0:01:00, bd
AssocID: dac2e123 in Active state, idle:0:00:18, timeout:0:01:00, os
```

- **show service-policy inspect m3ua table**

  Displays the run-time M3UA inspection table, including classification rules.

- **show conn detail**

  Displays connection information. M3UA connections are marked with the v flag.

History for Mobile Network Inspection

Feature Name	Releases	Feature Information
GTPv2 inspection and improvements to GTPv0/1 inspection.	9.5(1)	GTP inspection can now handle GTPv2. In addition, GTP inspection for all versions now supports IPv6 addresses. We changed the **GTP Inspect Map** > **Inspections** dialog box to let you configure separate message ID matching for GTPv1 and GTPv2. On the **General** parameters tab, the **GSN** timeout is now the **Endpoint** timeout.

Feature Name	Releases	Feature Information
SCTP inspection	9.5(2)	You can now apply application-layer inspection to Stream Control Transmission Protocol (SCTP) traffic to apply actions based on payload protocol identifier (PPID). We added or changed the following screens: **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **SCTP**; **Configuration** > **Firewall** > **Service Policy** add/edit wizard's **Rule Actions** > **Protocol Inspection** tab.
Diameter inspection	9.5(2)	You can now apply application-layer inspection to Diameter traffic and also apply actions based on application ID, command code, and attribute-value pair (AVP) filtering. We added or changed the following screens: **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **Diameter** and **Diameter AVP**; **Configuration** > **Firewall** > **Service Policy** add/edit wizard's **Rule Actions** > **Protocol Inspection** tab.
Diameter inspection improvements	9.6(1)	You can now inspect Diameter over TCP/TLS traffic, apply strict protocol conformance checking, and inspect Diameter over SCTP in cluster mode. We added or changed the following screens: **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **Diameter**; **Configuration** > **Firewall** > **Service Policy** add/edit wizard's **Rule Actions** > **Protocol Inspection** tab.
SCTP stateful inspection in cluster mode	9.6(1)	SCTP stateful inspection now works in cluster mode. You can also configure SCTP stateful inspection bypass in cluster mode. We did not add or modify any screens.
MTP3 User Adaptation (M3UA) inspection.	9.6(2)	You can now inspect M3UA traffic and also apply actions based on point code, service indicator, and message class and type. We added or modified the following pages: **Configuration** > **Firewall** > **Objects** > **Inspection Maps** > **M3UA**; the **Rule Action** > **Protocol Inspection** tab for service policy rules.
Support for SCTP multi-streaming reordering and reassembly and fragmentation. Support for SCTP multi-homing, where the SCTP endpoints have more than one IP address.	9.7(1)	The system now fully supports SCTP multi-streaming reordering, reassembly, and fragmentation, which improves Diameter and M3UA inspection effectiveness for SCTP traffic. The system also supports SCTP multi-homing, where the endpoints have more than one IP address each. For multi-homing, the system opens pinholes for the secondary addresses so that you do not need to write access rules to allow them. SCTP endpoints must be limited to 3 IP addresses each. We did not modify any ASDM screens.

Feature Name	Releases	Feature Information
M3UA inspection improvements.	9.7(1)	M3UA inspection now supports stateful failover, semi-distributed clustering, and multihoming. You can also configure strict application server process (ASP) state validation and validation for various messages. Strict ASP state validation is required for stateful failover and clustering. We modified the following screens: **Configuration** > **Firewall** > **Objects** > **Inspection Maps** > **M3UA** Add/Edit dialog boxes.
Support for setting the TLS proxy server SSL cipher suite.	9.8(1)	You can now set the SSL cipher suite when the ASA acts as a TLS proxy server. Formerly, you could only set global settings for the ASA on the **Configuration** > **Device Management** > **Advanced** > **SSL Settings** > **Encryption** page. We modified the following screen: **Configuration** > **Firewall** > **Unified Communications** > **TLS Proxy**, Add/Edit dialog boxes, **Server Configuration** page.
GTP inspection enhancements for MSISDN and Selection Mode filtering, anti-replay, and user spoofing protection.	9.10(1)	You can now configure GTP inspection to drop Create PDP Context messages based on Mobile Station International Subscriber Directory Number (MSISDN) or Selection Mode. You can also implement anti-replay and user spoofing protection. We modified the **Configuration** > **Firewall** > **Objects** > **Inspection Maps** > **GTP** > **Add/Edit** dialog box.
GTPv1 release 10.12 support.	9.12(1)	The system now supports GTPv1 release 10.12. Previously, the system supported release 6.1. The new support includes recognition of 25 additional GTPv1 messages and 66 information elements. In addition, there is a behavior change. Now, any unknown message IDs are allowed. Previously, unknown messages were dropped and logged. We did not add or change any screens.
Location logging for mobile stations (GTP inspection).	9.13(1)	You can configure GTP inspection to log the initial location of a mobile station and subsequent changes to the location. Tracking location changes can help you identify possibly fraudulent roaming charges. We modified the following screen: **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **GTP**.
GTPv2 and GTPv1 release 15 support.	9.13(1)	The system now supports GTPv2 3GPP 29.274 V15.5.0. For GTPv1, support is up to 3GPP 29.060 V15.2.0. The new support includes recognition of 2 additional messages and 53 information elements. We did not add or change any screens.

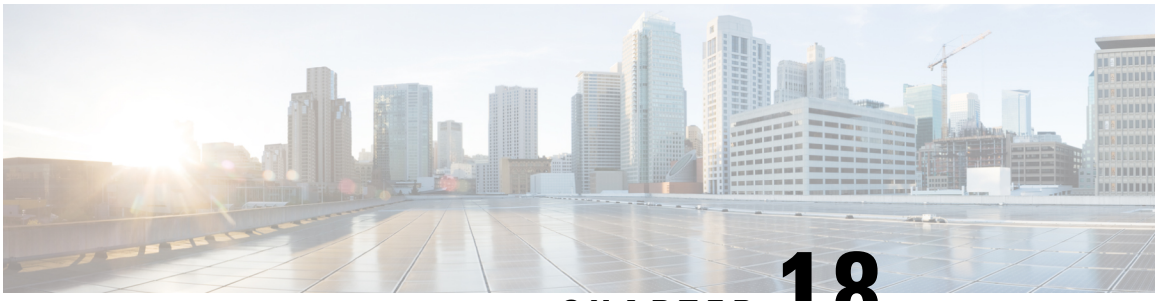
Feature Name	Releases	Feature Information
Ability to specify the IMSI prefixes to be dropped in GTP inspection.	9.16(1)	GTP inspection lets you configure IMSI prefix filtering, to identify the Mobile Country Code/Mobile Network Code (MCC/MNC) combinations to allow. You can now do IMSI filtering on the MCC/MNC combinations that you want to drop. This way, you can list out the unwanted combinations, and default to allowing all other combinations. We changed the following screens: The Drop option was added to the IMSI Prefix Filtering tab for GTP inspection maps.



PART V

Connection Management and Threat Detection

- Connection Settings, on page 413
- Quality of Service, on page 437
- Threat Detection, on page 447



CHAPTER 18

Connection Settings

This chapter describes how to configure connection settings for connections that go through the ASA, or for management connections that go to the ASA.

- Connection settings, on page 413
- Configure Connection Settings, on page 414
- Monitoring Connections, on page 432
- History for Connection Settings, on page 433

Connection settings

Connection settings refer to a variety of features that manage connections through a device.

Connection setting types

Connection settings include these types:

- **Global timeouts for various protocols**—All global timeouts have default values. Change them only if you experience premature connection loss.
- **Connection timeouts per traffic class**—You can override the global timeouts for specific types of traffic using service policies. All traffic class timeouts have default values, so you do not have to set them.
- **Connection limits and TCP Intercept**—By default, there are no limits on how many connections can go through (or to) the ASA. Use service policy rules to set limits on particular traffic classes. This protects servers from denial of service attacks Particularly, you can set limits on embryonic connections (those that have not finished the TCP handshake), which protects against SYN flooding attacks. When embryonic limits are exceeded, the TCP Intercept component activates to proxy connections and throttle attacks.
- **Dead connection detection (DCD)**—Enable dead connection detection if you have persistent connections that are valid but often idle. These connections may get closed when they exceed idle timeout settings. Dead connection detection identifies idle but valid connections and keeps them alive by resetting their idle timers. Whenever idle times are exceeded, DCD probes both sides of the connection to see if both sides agree that the connection is valid.. The **show service-policy** command output includes counters to show the amount of activity from DCD. You can use the **show conn detail** command to get information about the initiator and responder and how often each has sent probes.
- **TCP sequence randomization**—Each TCP connection has two initial sequence numbers (ISN): one generated by the client and one generated by the server. By default, the ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions. Randomization prevents an attacker

from predicting the next ISN for a new connection and potentially hijacking the new session. However, TCP sequence randomization effectively breaks TCP SACK (Selective Acknowledgement), as the sequence numbers the client sees are different from what the server sees. You can disable randomization per traffic class if desired.

- **TCP normalization**—The TCP Normalizer protects against abnormal packets. You can configure how some types of packet abnormalities are handled by traffic class.
- **TCP state bypass**—You can bypass TCP state checking if you use asymmetrical routing in your network.
- **SCTP state bypass**—You can bypass Stream Control Transmission Protocol (SCTP) stateful inspection if you do not want SCTP protocol validation.
- **Flow offloading**—You can identify select traffic to be offloaded to a super fast path, where the flows are switched in the NIC itself. Offloading can help you improve performance for data-intensive applications such as large file transfers.

Configure Connection Settings

Connection limits, timeouts, TCP Normalization, TCP sequence randomization, and decrementing time-to-live (TTL) have default values that are appropriate for most networks. You need to configure these connection settings only if you have unusual requirements, your network has specific types of configuration, or if you are experiencing unusual connection loss due to premature idle timeouts.

Other connection-related features are not enabled. You would configure these services on specific traffic classes only, and not as a general service. These features include the following: TCP Intercept, TCP State Bypass, Dead Connection Detection (DCD), SCTP state bypass, flow offload.

The following general procedure covers the gamut of possible connection setting configurations. Pick and choose which to implement based on your needs.

Before you begin

When a packet enters the device, the firewall first evaluates access control rules and NAT to determine whether a connection entry needs to be created. A connection entry is created regardless of whether an ARP entry exists for the next hop. Thus, the existence of a connection does not mean that any packet within the scope of the connection made it through the device. Connection settings such as idle timeouts ensure that unwanted connections do not linger and take up system resources.

Procedure

Step 1 Configure Global Timeouts, on page 415. These settings change the default idle timeouts for various protocols for all traffic that passes through the device. If you are having problems with connections being reset due to premature timeouts, first try changing the global timeouts.

Step 2 Protect Servers from a SYN Flood DoS Attack (TCP Intercept), on page 417. Use this procedure to configure TCP Intercept.

Step 3 Customize Abnormal TCP Packet Handling (TCP Maps, TCP Normalizer), on page 419, if you want to alter the default TCP Normalization behavior for specific traffic classes.

Step 4 Bypass TCP state checks for asymmetrical routing (TCP state bypass), on page 422, if you have this type of routing environment.

Step 5 Disable TCP Sequence Randomization, on page 424, if the default randomization is scrambling data for certain connections.

Step 6 Offload Large Flows, on page 425, if you need to improve performance in a computing intensive data center.

Step 7 Configure Connection Settings for Specific Traffic Classes (All Services), on page 428. This is a catch-all procedure for connection settings. These settings can override the global defaults for specific traffic classes using service policy rules. You also use these rules to customize TCP Normalizer, change TCP sequence randomization, decrement time-to-live on packets, and implement other optional features.

Step 8 Configure TCP Options, on page 431, if you need to force resets or change some other standard TCP behavior.

Configure Global Timeouts

You can set the global idle timeout durations for the connection and translation slots of various protocols. If the slot has not been used for the idle time specified, the resource is returned to the free pool.

Changing the global timeout sets a new default timeout, which in some cases can be overridden for particular traffic flows through service policies.

For protocols that do not have a specific configurable timeout setting, such as GRE, the idle timeout is 2 minutes.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Advanced** > **Global Timeouts**.

Step 2 Configure the timeouts by checking the boxes for timeouts you want to change and entering the new value.

All durations are displayed in the format *hh*:*mm*:*ss*, with a maximum duration of 1193:0:0 in most cases. In all cases, except for Authentication absolute and Authentication inactivity, unchecking the check boxes returns the timeout to the default value. For those two cases, clearing the check box means to reauthenticate on every new connection.

Enter 0 to disable a timeout.

- **Connection**—The idle time until a connection slot is freed. This duration must be at least 5 minutes. The default is 1 hour.
- **Half-closed**—The idle time until a TCP half-closed connection closes. A connection is considered half-closed if both the FIN and FIN-ACK have been seen. If only the FIN has been seen, the regular connection timeout applies. The minimum is 30 seconds. The default is 10 minutes.
- **UDP**—The idle time until a UDP connection closes. This duration must be at least 1 minute. The default is 2 minutes.
- **ICMP**—The idle time after which general ICMP states are closed. The default (and minimum) is 2 seconds.
- **ICMP Error**—The idle time before the ASA removes an ICMP connection after receiving an ICMP echo-reply packet, between 0:0:0 and 0:1:0 or the ICMP timeout value, whichever is lower. The default is 0 (disabled). When this timeout is disabled, and you enable ICMP inspection, then the ASA removes

the ICMP connection as soon as an echo-reply is received; thus any ICMP errors that are generated for the (now closed) connection are dropped. This timeout delays the removal of ICMP connections so you can receive important ICMP errors.

- **H.323**—The idle time after which H.245 (TCP) and H.323 (UDP) media connections close. The default (and minimum) is 5 minutes. Because the same connection flag is set on both H.245 and H.323 media connections, the H.245 (TCP) connection shares the idle timeout with the H.323 (RTP and RTCP) media connection.
- **H.225**—The idle time until an H.225 signaling connection closes. The default is 1 hour. To close a connection immediately after all calls are cleared, a timeout of 1 second (0:0:1) is recommended.
- **MGCP**—The idle time after which an MGCP media connection is removed. The default is 5 minutes, but you can set it as low as 1 second.
- **MGCP PAT**—The idle time after which an MGCP PAT translation is removed. The default is 5 minutes. The minimum time is 30 seconds.
- **TCP Proxy Reassembly**—The idle timeout after which buffered packets waiting for reassembly are dropped, between 0:0:10 and 1193:0:0. The default is 1 minute (0:1:0).
- **Floating Connection**—When multiple routes exist to a network with different metrics, the system uses the one with the best metric at the time of connection creation. If a better route becomes available, then this timeout lets connections be closed so a connection can be reestablished to use the better route. The default is 0 (the connection never times out). To make it possible to use better routes, set the timeout to a value between 0:0:30 and 1193:0:0. This timer does not apply to connections through virtual tunnel interfaces (VTI). If a connection through a VTI gets stuck, you must manually clear it.
- **SCTP**—The idle time until a Stream Control Transmission Protocol (SCTP) connection closes, between 0:1:0 and 1193:0:0. The default is 2 minutes (0:2:0).
- **Stale Routes**—How long to keep a stale route before removing it from the router information base. These routes are for interior gateway protocols such as OSPF. The default is 70 seconds (00:01:10), the range is 00:00:10 to 00:01:40.
- **SUNRPC**—The idle time until a SunRPC slot is freed. This duration must be at least 1 minute. The default is 10 minutes.
- **SIP**—The idle time until a SIP signaling port connection closes. This duration must be at least 5 minutes. The default is 30 minutes.
- **SIP Media**—The idle time until a SIP media port connection closes. This duration must be at least 1 minute. The default is 2 minutes. The SIP media timer is used for SIP RTP/RTCP with SIP UDP media packets, instead of the UDP inactivity timeout.
- **SIP Provisional Media**—The timeout value for SIP provisional media connections, between 1 and 30 minutes. The default is 2 minutes.
- **SIP Invite**—The idle time after which pinholes for PROVISIONAL responses and media xlates will be closed, between 0:1:0 and 00:30:0. The default is 3 minutes (0:3:0).
- **SIP Disconnect**—The idle time after which SIP session is deleted if the 200 OK is not received for a CANCEL or a BYE message, between 0:0:1 and 0:10:0. The default is 2 minutes (0:2:0).
- **Authentication absolute**—The duration until the authentication cache times out and users have to reauthenticate a new connection. This timer is used in cut-through proxy only, which is a AAA rule. This duration must be shorter than the Translation Slot timeout. The system waits until the user starts a new

connection to prompt again. Before you disable caching to force authentication on every new connection, consider the following limitations.

- Do not set this value to 0 if passive FTP is used on the connections.
- When Authentication Absolute is 0, HTTPS authentication may not work. If a browser initiates multiple TCP connections to load a web page after HTTPS authentication, the first connection is permitted through, but subsequent connections trigger authentication. As a result, users are continuously presented with an authentication page, even after successful authentication. To work around this, set the authentication absolute timeout to 1 second. This workaround opens a 1-second window of opportunity that might allow non-authenticated users to go through the firewall if they are coming from the same source IP address.

- **Authentication inactivity**—The idle time until the authentication cache times out and users have to reauthenticate a new connection. This duration must be shorter than the Translation Slot value. This timeout is disabled by default. This timer is used in cut-through proxy only, which is a AAA rule.
- **Translation Slot**—The idle time until a NAT translation slot is freed. This duration must be at least 1 minute. The default is 3 hours.
- **PAT Translation Slot** (8.4(3) and later, not including 8.5(1) and 8.6(1))—The idle time until a PAT translation slot is freed, between 0:0:30 and 0:5:0. The default is 30 seconds. You may want to increase the timeout if upstream routers reject new connections using a freed PAT port because the previous connection might still be open on the upstream device.
- **Connection Holddown**—How long the system should maintain a connection when the route used by the connection no longer exists or is inactive. If the route does not become active within this holddown period, the connection is freed. The purpose of the connection holddown timer is to reduce the effect of route flapping, where routes might come up and go down quickly. You can reduce the holddown timer to make route convergence happen more quickly. The default is 15 seconds, the range is 00:00:00 to 00:00:15.

Step 3 Click **Apply**.

Protect Servers from a SYN Flood DoS Attack (TCP Intercept)

A SYN-flooding denial of service (DoS) attack occurs when an attacker sends a series of SYN packets to a host. These packets usually originate from spoofed IP addresses. The constant flood of SYN packets keeps the server SYN queue full, which prevents it from servicing connection requests from legitimate users.

You can limit the number of embryonic connections to help prevent SYN flooding attacks. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination.

When the embryonic connection threshold of a connection is crossed, the ASA acts as a proxy for the server and generates a SYN-ACK response to the client SYN request using the SYN cookie method, so that the connection is not added to the SYN queue of the targeted host. The SYN cookie is the initial sequence number returned in the SYN-ACK that is constructed from MSS, time stamp, and a mathematical hash of other items to essentially create a secret. If the ASA receives an ACK back from the client with the correct sequence number and within the valid time window, it can then authenticate that the client is real and allow the connection to the server. The component that performs the proxy is called TCP Intercept.

The end-to-end process for protecting a server from a SYN flood attack involves setting connection limits, enabling TCP Intercept statistics, and then monitoring the results.

Before you begin

- Ensure that you set the embryonic connection limit lower than the TCP SYN backlog queue on the server that you want to protect. Otherwise, valid clients can no longer access the server during a SYN attack. To determine reasonable values for embryonic limits, carefully analyze the capacity of the server, the network, and server usage.
- Depending on the number of CPU cores on your ASA model, the maximum concurrent and embryonic connections can exceed the configured numbers due to the way each core manages connections. In the worst case scenario, the ASA allows up to *n*-1 extra connections and embryonic connections, where *n* is the number of cores. For example, if your model has 4 cores, if you configure 6 concurrent connections and 4 embryonic connections, you could have an additional 3 of each type. To determine the number of cores for your model, enter the **show cpu core** command.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy**.

Step 2 Click **Add** > **Add Service Policy Rule**.

Alternatively, if you already have a rule for the servers you want to protect, edit the rule.

Step 3 Select whether to apply the rule to a specific interface or globally to all interfaces, and click **Next**.

Step 4 For Traffic Classification, select **Source and Destination IP Addresses (uses ACL)** and click **Next**.

Step 5 For the ACL rule, enter the IP addresses of the servers in **Destination**, and specify the protocol for the servers. Typically, you would use **any** for the **Source**. Click **Next** when finished.

For example, if you want to protect the web servers 10.1.1.5 and 10.1.1.6, enter:

- Source = any
- Destination = 10.1.1.5, 10.1.1.6
- Destination Protocol = tcp/http

Step 6 On the Rule Actions page, click the **Connection Settings** tab and fill in these options:

- **Embryonic Connections**—The maximum number of embryonic TCP connections per host up to 2000000. The default is **0**, which means the maximum embryonic connections are allowed. For example, you could set this to 1000.
- **Per Client Embryonic Connections**—The maximum number of simultaneous embryonic TCP connections for each client up to 2000000. When a new TCP connection is requested by a client that already has the maximum per-client number of embryonic connections open through the ASA, the ASA prevents the connection. For example, you could set this to 50.
- **TCP Syn Cookie MSS**—The server maximum segment size (MSS) for SYN-cookie generation for embryonic connections upon reaching the embryonic connections limit, from 48 to 65535 . The default is 1380. This setting is meaningful only if you configure **Embryonic Connections** or **Per Client Embryonic Connections**.

Step 7 Click **Finish** to save the rule, and **Apply** to update the device.

Step 8 Choose **Configuration** > **Firewall** > **Threat Detection**, and enable at least the **TCP Intercept** statistics under the Threat Detection Statistics group.

You can simply enable all statistics, or just enable TCP Intercept. You can also adjust the monitoring window and rates.

Step 9 Choose **Home** > **Firewall Dashboard**, and look at the **Top Ten Protected Servers under SYN Attack** dashboard to monitor the results.

Click the **Detail** button to show history sampling data. The ASA samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

You can clear the statistics by entering the **clear threat-detection statistics tcp-intercept** command using **Tools > Command Line Interface.**

Customize Abnormal TCP Packet Handling (TCP Maps, TCP Normalizer)

The TCP normalizer identifies abnormal packets that the ASA can act on when they are detected; for example, the ASA can allow, drop, or clear the packets. TCP normalization helps protect the ASA from attacks. TCP normalization is always enabled, but you can customize how some features behave.

To customize the TCP normalizer, first define the settings using a TCP map. Then, you can apply the map to selected traffic classes using service policies.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **TCP Maps**.

Step 2 Do one of the following:

- Click **Add** to add a new TCP map. Enter a name for the map.
- Select a map and click **Edit**.

Step 3 In the **Queue Limit** field, enter the maximum number of out-of-order packets that can be buffered and put in order for a TCP connection, between 0 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:

- Connections for application inspection and TCP check-retransmission have a queue limit of 3 packets. If the ASA receives a TCP packet with a different window size, then the queue limit is dynamically changed to match the advertised setting.
- For other TCP connections, out-of-order packets are passed through untouched.

If you set the Queue Limit to be 1 or above, then the number of out-of-order packets allowed for all TCP traffic matches this setting. For example, for application inspection and TCP check-retransmission traffic, any advertised settings from TCP packets are ignored in favor of the Queue Limit setting. For other TCP traffic, out-of-order packets are now buffered and put in order instead of passed through untouched.

Step 4 In the **Timeout** field, set the maximum amount of time that out-of-order packets can remain in the buffer, between 1 and 20 seconds.

If they are not put in order and passed on within the timeout period, then they are dropped. The default is 4 seconds. You cannot change the timeout for any traffic if the Queue Limit is set to 0; you need to set the limit to be 1 or above for the Timeout to take effect.

Step 5 For **Reserved Bits**, select how to handle packets that have reserved bits in the TCP header: **Clear and allow** (remove the bits before allowing the packet), **Allow only** (do not change the bits, the default), or **Drop** the packet.

Step 6 Select any of the following options:

- **Clear urgent flag**—Clears the URG flag in a packet before allowing it. The URG flag is used to indicate that the packet contains information that is of higher priority than other data within the stream. The TCP RFC is vague about the exact interpretation of the URG flag, therefore end systems handle urgent offsets in different ways, which may make the end system vulnerable to attacks.
- **Drop connection on window variation**—Drops a connection that has changed its window size unexpectedly. The window size mechanism allows TCP to advertise a large window and to subsequently advertise a much smaller window without having accepted too much data. From the TCP specification, "shrinking the window" is strongly discouraged.
- **Drop packets that exceed maximum segment size**—Drops packets that exceed the MSS set by the peer.
- **Check if transmitted data is the same as original**—Enables the retransmit data checks, which prevent inconsistent TCP retransmissions.
- **Drop packets which have past-window sequence**—Drops packets that have past-window sequence numbers, namely the sequence number of a received TCP packet is greater than the right edge of the TCP receiving window. To allow these packets, deselect this option and set the **Queue Limit** to 0 (disabling the queue limit).
- **Drop SYN Packets with data**—Drops SYN packets that contain data.
- **Enable TTL Evasion Protection**—Have the maximum TTL for a connection be determined by the TTL in the initial packet. The TTL for subsequent packets can decrease, but it cannot increase. The system will reset the TTL to the lowest previously-seen TTL for that connection. This protects against TTL evasion attacks.

  For example, an attacker can send a packet that passes policy with a very short TTL. When the TTL goes to zero, a router between the ASA and the endpoint drops the packet. It is at this point that the attacker can send a malicious packet with a long TTL that appears to the ASA to be a retransmission and is passed. To the endpoint host, however, it is the first packet that has been received by the attacker. In this case, an attacker is able to succeed without security preventing the attack.
- **Verify TCP Checksum**—Verifies the TCP checksum, dropping packets that fail verification.
- **Drop SYNACK Packets with data**—Drops TCP SYNACK packets that contain data.
- **Drop packets with invalid ACK**—Drops packets with an invalid ACK. You might see invalid ACKs in the following instances:
  - In the TCP connection SYN-ACK-received status, if the ACK number of a received TCP packet is not exactly same as the sequence number of the next TCP packet sending out, it is an invalid ACK.

- Whenever the ACK number of a received TCP packet is greater than the sequence number of the next TCP packet sending out, it is an invalid ACK.

Note
TCP packets with an invalid ACK are automatically allowed for WAAS connections.

Step 7 (Optional.) Click the **TCP Options** tab and set the action for packets that contain TCP options.

You can clear the options before allowing the packets, allow the packets if they contain a single option of a given type, or allow the packets even if they have more than one option of a given type. The default is to allow the five named options as long as a given option appears no more than once per packet (otherwise the packet is dropped), while clearing all other options. You can also elect to drop packets that contain the MD5 or any of the numbered options. Note that if a TCP connection is inspected, all options are cleared except the MSS and selective-acknowledgment (SACK) options, regardless of your configuration.

a) Select the action for the **Selective Acknowledgement**, **TCP Timestamp**, and **Window Scale** options.

Clearing the timestamp option disables PAWS and RTT.

b) Select the action for the **MSS** (Maximum Segment Size) option.

In addition to the regular allow, allow multiple, and clear actions, you can select **Specify Maximum** and enter the maximum segment size, from 68-65535. The default TCP MSS is defined on the **Configuration** > **Firewall** > **Advanced** > **TCP Options** page.

c) Select whether you want to **Allow packets with the MD5 option**.

If you deselect the checkbox, packets that contain the MD5 option are dropped. If you select the option, you can apply the normal actions of allow, allow multiple, or clear.

d) Select the action for options by number range.

Options numbered 6-7, 9-18, and 20-255 are cleared by default. You can instead allow the options, or drop packets that contain the option. You can specify different actions for different option ranges: simply enter the lower and upper number for the range, select the action, and click **Add**. To configure an action for a single option, enter the same number for the lower and upper range.

To remove a configured range, select it and click **Delete**.

Step 8 Click **OK** and **Apply**.

You can now use the TCP map in a service policy. The map affects traffic only when applied through a service policy.

Step 9 Apply the TCP map to a traffic class using a service policy.

a) Choose **Configuration > Firewall > Service Policy Rules**.
b) Add or edit a rule. You can apply the rule globally or to an interface. For example, to customize abnormal packet handling for all traffic, create a global rule that matches any traffic. Proceed to the Rule Actions page.
c) Click the **Connection Settings** tab.
d) Choose **Use TCP Map** and select the map you created.
e) Click **Finish** or **OK**, then click **Apply**.

Bypass TCP state checks for asymmetrical routing (TCP state bypass)

TCP state bypass allows traffic to bypass stateful inspection checks on ASA devices.

Implementation considerations

If you have an asymmetrical routing environment in your network, where the outbound and inbound flow for a given connection can go through two different ASA devices, you need to implement TCP state bypass on the affected traffic.

However, TCP state bypass weakens the security of your network, so you should apply bypass on very specific, limited traffic classes.

The asymmetrical routing problem

The asymmetrical routing problem occurs when return traffic for a connection goes through a different firewall than the outbound traffic. This causes packets to be dropped because they do not match existing connection entries in the fast path table.

Traffic inspection and fast path operation

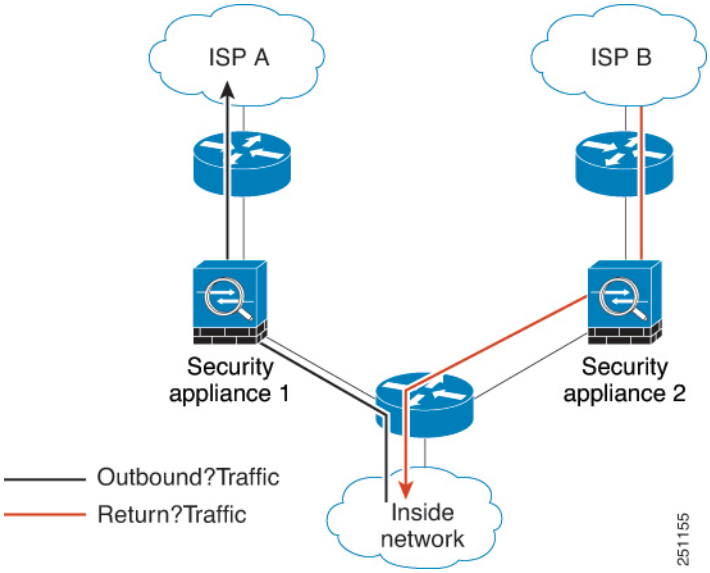
By default, the ASA inspects all traffic using the Adaptive Security Algorithm. The ASA allows or drops traffic based on the security policy. The ASA maximizes the firewall performance by checking the state of each packet (new connection or established connection). It assigns new connection SYN packets to the session management path. It assigns established connections to the fast path. It assigns packets requiring advanced inspection to the control plane path.

The ASA allows TCP packets that match existing connections to pass through via the fast path without rechecking every aspect of the security policy. This feature improves performance. However, the fast path establishment method can interfere with asymmetric routing solutions. The method uses SYN packets and performs checks such as TCP sequence number validation. Both outbound and inbound flows of a connection must pass through the same ASA device.

Asymmetric routing scenario

For example, a new connection goes to Security Appliance 1. The SYN packet goes through the session management path, and an entry for the connection is added to the fast path table. If subsequent packets of this connection go through Security Appliance 1, then the packets match the entry in the fast path, and are passed through. But if subsequent packets go to Security Appliance 2, where no SYN packet went through the session management path, then no entry exists in the fast path for the connection, and Security Appliance 2 drops the packets. The following figure shows an asymmetric routing example where the outbound traffic goes through a different ASA than the inbound traffic:

Figure 50: Asymmetric routing



If you have asymmetric routing configured on upstream routers and traffic alternates between two ASA devices, then you can configure TCP state bypass for specific traffic. TCP state bypass alters the way sessions are established in the fast path and disables the fast path checks. This feature treats TCP traffic similarly to a UDP connection: when a non-SYN packet matching the specified networks enters the ASA device, and there is not a fast path entry, then the packet goes through the session management path to establish the connection in the fast path. Once in the fast path, the traffic bypasses the fast path checks.

Guidelines and limitations for TCP state bypass

TCP state bypass unsupported features

The system does not support these features when you use TCP state bypass:

- Application inspection—Inspection requires both inbound and outbound traffic to go through the same ASA, so inspection is not applied to TCP state bypass traffic.
- AAA authenticated sessions—When a user authenticates with one ASA, traffic returning via the other ASA will be denied because the user did not authenticate with that ASA.
- TCP Intercept, maximum embryonic connection limit, TCP sequence number randomization—The ASA does not keep track of the state of the connection, so these features are not applied.
- TCP normalization—The TCP normalizer is disabled.
- Service module functionality—You cannot use TCP state bypass and any application running on any type of service module, such as ASA FirePOWER.

TCP state bypass NAT guidelines

Because the translation session is established separately for each device, be sure to configure static NAT on both devices for TCP state bypass traffic. If you use dynamic NAT, the address chosen for the session on Device 1 will differ from the address chosen for the session on Device 2.

Configure TCP State Bypass

To bypass TCP state checking in asymetrical routing environments, carefully define a traffic class that applies to the affected hosts or networks only, then enable TCP State Bypass on the traffic class using a service policy. Because bypass reduces the security of the network, limit its application as much as possible.

Before you begin

If there is no traffic on a given connection for 2 minutes, the connection times out. You can override this default by changing the **Idle Connection Timeout** for the TCP state bypass traffic class. Normal TCP connections timeout by default after 60 minutes.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy**.

Step 2 Click **Add** > **Add Service Policy Rule**.

Alternatively, if you already have a rule for the hosts, edit the rule.

Step 3 Select whether to apply the rule to a specific interface or globally to all interfaces, and click **Next**.

Step 4 For Traffic Classification, select **Source and Destination IP Addresses (uses ACL)** and click **Next**.

Step 5 For the ACL rule, enter the IP addresses of the hosts on each end of the route in **Source** and **Destination**, and specify the protocol as TCP. Click **Next** when finished.

For example, if you want to bypass TCP state checking between 10.1.1.1 and 10.2.2.2, enter:

- Source = 10.1.1.1
- Destination = 10.2.2.2
- Destination Protocol = tcp

Step 6 On the Rule Actions page, click the **Connection Settings** tab and select **TCP State Bypass**.

Step 7 Click **Finish** to save the rule, and **Apply** to update the device.

Disable TCP Sequence Randomization

Each TCP connection has two ISNs: one generated by the client and one generated by the server. The ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session. However, TCP sequence randomization effectively breaks TCP SACK (Selective Acknowledgement), as the sequence numbers the client sees are different from what the server sees.

You can disable TCP initial sequence number randomization if necessary, for example, because data is getting scrambled. For example:

- If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.

- If you use eBGP multi-hop through the ASA, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
- You use a WAAS device that requires the ASA not to randomize the sequence numbers of connections.
- You enable hardware bypass for the ISA 3000, and TCP connections are dropped when the ISA 3000 is no longer part of the data path.



Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy**.

Step 2 Click **Add** > **Add Service Policy Rule**.

Alternatively, if you already have a rule for the targeted traffic, edit the rule.

Step 3 Select whether to apply the rule to a specific interface or globally to all interfaces, and click **Next**.

Step 4 For Traffic Classification, identity the type of traffic match. The class match should be for TCP traffic; you can identify specific hosts (with an ACL) do a TCP port match, or simply match any traffic. Click **Next** and configure the hosts in the ACL or define the ports, and click **Next** again.

For example, if you want to disable TCP sequence number randomization for all TCP traffic directed at 10.2.2.2, enter:

- Source = any
- Destination = 10.2.2.2
- Destination Protocol = tcp

Step 5 On the Rule Actions page, click the **Connection Settings** tab and uncheck **Randomize Sequence Number**.

Step 6 Click **Finish** to save the rule, and **Apply** to update the device.

Offload Large Flows

If you deploy the ASA on supported devices in a data center, you can identify select traffic to be offloaded to a super fast path, where traffic is switched in the NIC itself. Offloading can help you improve performance for data-intensive applications such as large file transfers.

- High Performance Computing (HPC) Research sites, where the ASA is deployed between storage and high compute stations. When one research site backs up using FTP file transfer or file sync over NFS, the large amount of data traffic affects all contexts on the ASA. Offloading FTP file transfer and file sync over NFS reduces the impact on other traffic.

- High Frequency Trading (HFT), where the ASA is deployed between workstations and the Exchange, mainly for compliance purposes. Security is usually not a concern, but latency is a major concern.

Before being offloaded, the ASA first applies normal security processing, such as access rules and inspection, during connection establishment. The ASA also does session tear-down. But once a connection is established, if it is eligible to be offloaded, further processing happens in the NIC rather than the ASA.

Offloaded flows continue to receive limited stateful inspection, such as basic TCP flag and option checking, and checksum verification if you configure it. The system can selectively escalate packets to the firewall system for further processing if necessary.

To identify flows that can be offloaded, you create a service policy rule that applies the flow offloading service. A matching flow is then offloaded if it meets the following conditions:

- IPv4 addresses only.
- TCP, UDP, GRE only.
- Standard or 802.1Q tagged Ethernet frames only.
- (Transparent mode only.) Multicast flows for bridge groups that contain two and only two interfaces.

Reverse flows for offloaded flows are also offloaded.

Restrictions for flow offload implementation

Device limitations

The feature is supported on the following devices:

- Firepower 4100/9300

Flows that cannot be offloaded

Do not attempt to offload these flow types:

- Any flows that do not use IPv4 addressing, such as IPv6 addressing.
- Flows for any protocol other than TCP, UDP, and GRE.



Note PPTP GRE connections cannot be offloaded.

- Flows that require inspection. In some cases, such as FTP, the secondary data channel can be offloaded although the control channel cannot be offloaded.
- Flows that pass through another module, such as ASA FirePOWER.
- IPsec and TLS/DTLS VPN connections that terminate on the device.
- Multicast flows in routed mode.
- Multicast flows in transparent mode for bridge groups that have three or more interfaces.
- TCP Intercept flows.

- TCP state bypass flows. You cannot configure flow offload and TCP state bypass on the same traffic.
- AAA cut-through proxy flows.
- Vpath, VXLAN related flows.
- Flows tagged with security groups.
- Reverse flows that are forwarded from a different cluster node, in the case of asymmetric flows in a cluster.
- Centralized flows in a cluster, if the flow owner is not the control unit.

Additional limitations

Consider these operational restrictions when implementing flow offload:

- Flow offload and Dead Connection Detection (DCD) are not compatible. Do not configure DCD on connections that can be offloaded.
- If more than one flow that matches flow offload conditions are queued to be offloaded at the same time to the same location on the hardware, only the first flow is offloaded. The other flows are processed normally. This is called a *collision*. Use the **show flow-offload flow** command in the CLI to display statistics for this situation.
- Although offloaded flows pass through FXOS interfaces, statistics for these flows do not appear on the logical device interface. Thus, logical device interface counters and packet rates do not reflect offloaded flows.

Conditions for reversing offload

After a flow is offloaded, packets within the flow are returned to the firewall operating system for further processing if they meet these conditions:

- They include TCP options other than Timestamp.
- They are fragmented.
- They are subject to Equal-Cost Multi-Path (ECMP) routing, and ingress packets move from one interface to another.
- A change to the routing table applies to the flow. A packet is returned to the device to determine the new route.

Configure Flow Offload

To configure flow offload, you must enable the service and then create service policies to identify the traffic that is eligible for offloading.When you initially enable the service, you must reboot. .

Procedure

Step 1 Enable the flow offload service.

When you initially enable the service, you must reboot.

If a reload is required, there are special considerations for clusters or failover pairs if you want a hitless change:

- Clustering—First enter the command on the control node, but do not reboot the control node immediately. Instead, reboot each node of the cluster first, then return to the control node and reboot it. You can then configure the offloading service policy on the control node.
- Failover—First enter the command on the active unit, but do not reboot it immediately. Instead, reboot the standby unit, then reboot the active unit. You can then configure the offloading service policy on the active unit.

a) Select **Configuration** > **Firewall** > **Advanced** > **Offload Engine**.
b) Check **Enable Offload Engine**.
c) Click **Apply**.
d) Click **Save** to save your changes to the startup configuration.
e) Select **Tools** > **System Reload** to reboot the device.

Step 2 Create the service policy rule that identifies traffic that is eligible for offload.

a) Choose **Configuration** > **Firewall** > **Service Policy**.
b) Click **Add** > **Add Service Policy Rule**.

Alternatively, if you already have a rule for the hosts, edit the rule.

c) Select whether to apply the rule to a specific interface or globally to all interfaces, and click **Next**.
d) For Traffic Classification, matching by access-list (**Source and Destination IP Addresses (uses ACL)**) or port (**TCP or UDP or SCTP Destination Port**) would be the most typical options. Select an option and click **Next**.
e) Enter the ACL or port criteria. Click **Next** when finished.

For example, if you want to make all TCP traffic on the 10.1.1.0/255.255.255.224 subnet eligible for offload, enter:

- Source = 10.1.1.0/255.255.255.224 (or 10.1.1.0/27)
- Destination = any
- Destination Protocol = tcp

f) On the Rule Actions page, click the **Connection Settings** tab and select **Flow Offload**.
g) Click **Finish** to save the rule, and **Apply** to update the device.

Configure Connection Settings for Specific Traffic Classes (All Services)

You can configure different connection settings for specific traffic classes using service policies. Use service policies to:

- Customize connection limits and timeouts used to protect against DoS and SYN-flooding attacks.
- Implement Dead Connection Detection so that valid but idle connections remain alive.
- Disable TCP sequence number randomization in cases where you do not need it.
- Customize how the TCP Normalizer protects against abnormal TCP packets.

- Implement TCP State Bypass for traffic subject to asymetrical routing. Bypass traffic is not subject to inspection.
- Implement Stream Control Transmission Protocol (SCTP) State Bypass to turn off SCTP stateful inspection.
- Implement flow offload to improve performance on supported hardware platforms.
- Decrement time-to-live (TTL) on packets so that the ASA will show up on trace route output.



Note If you decrement time to live, packets with a TTL of 1 will be dropped, but a connection will be opened for the session on the assumption that the connection might contain packets with a greater TTL. Note that some packets, such as OSPF hello packets, are sent with TTL = 1, so decrementing time to live can have unexpected consequences for transparent mode ASA devices. The decrement time-to-live settings does not impact the OSPF process when ASA is operating in a routed mode.

You can configure any combination of these settings for a given traffic class, except for TCP State Bypass and TCP Normalizer customization, which are mutually exclusive.



Tip This procedure shows a service policy for traffic that goes through the ASA. You can also configure the connection maximum and embryonic connection maximum for management (to the box) traffic.

Before you begin

If you want to customize the TCP Normalizer, create the required TCP Map before proceeding.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy**, and open a rule.

- To create a new rule, click **Add > Add Service Policy Rule**. Proceed through the wizard to the Rules page.
- If you have a rule for which you are changing connection settings, select it and click **Edit**.

Step 2 On the Rule Actions wizard page or tab, select the **Connection Settings** tab.

Step 3 To set maximum connections, configure the following values in the Maximum Connections area:

By default, there are no connection limits. If you implement limits, the system must start tracking them, which can increase CPU and memory usage and result in operational problems for systems under heavy load, especially in a cluster.

- **Maximum TCP, UDP and SCTP Connections**—(TCP, UDP, SCTP.) The maximum number of simultaneous connections for all clients in the traffic class, up to 2000000. The default is 0, which means

the maximum possible connections are allowed. For TCP connections, this applies to established connections only.

- **Embryonic Connections**—Specifies the maximum number of embryonic TCP connections per host up to 2000000. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. The default is **0**, which means the maximum embryonic connections are allowed. By setting a non-zero limit, you enable TCP Intercept, which protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. Also set the per-client options to protect against SYN flooding.
- **Per Client Connections**—(TCP, UDP, SCTP.) Specifies the maximum number of simultaneous connections for each client up to 2000000. When a new connection is attempted by a client that already has opened the maximum per-client number of connections, the ASA rejects the connection and drops the packet. For TCP connections, this includes established, half-open, and half-closed connections.
- **Per Client Embryonic Connections**—Specifies the maximum number of simultaneous TCP embryonic connections for each client up to 2000000. When a new TCP connection is requested by a client that already has the maximum per-client number of embryonic connections open through the ASA, the ASA prevents the connection.
- **TCP Syn Cookie MSS**—The server maximum segment size (MSS) for SYN-cookie generation for embryonic connections upon reaching the embryonic connections limit, from 48 to 65535 . The default is 1380. This setting is meaningful only if you configure **Embryonic Connections** or **Per Client Embryonic Connections**.

Step 4 To configure connection timeouts, configure the following values in the TCP Timeout area:

- **Embryonic Connection Timeout**—The idle time until an embryonic (half-open) TCP connection slot is freed. Enter 0:0:0 to disable timeout for the connection. The default is 30 seconds.
- **Half Closed Connection Timeout**—The idle timeout period until a half-closed connection is closed, between 0:5:0 (for 9.1(1) and earlier) or 0:0:30 (for 9.1(2) and later) and 1193:0:0. The default is 0:10:0. Half-closed connections are not affected by DCD. Also, the ASA does not send a reset when taking down half-closed connections.
- **Idle Connection Timeout**—The idle time until a connection slot (of *any* protocol, not just TCP) is freed. Enter 0:0:0 to disable timeout for the connection. This duration must be at least 5 minutes. The default is 1 hour.
- **Send reset to TCP endpoints before timeout**—Whether the ASA should send a TCP reset message to the endpoints of the connection before freeing the connection slot.
- **Dead Connection Detection (DCD)**—Whether to enable Dead Connection Detection (DCD). Before expiring an idle connection, the ASA probes the end hosts to determine if the connection is valid. If both hosts respond, the connection is preserved, otherwise the connection is freed. Set the maximum number of retries (default is 5, the range is 1-255) and the retry interval, which is the period to wait after each unresponsive DCD probe before sending another probe (0:0:1 to 24:0:0, default is 0:0:15). When operating in transparent firewall mode, you must configure static routes for the endpoints. You cannot configure DCD on connections that are also offloaded, so ensure DCD and flow offload traffic classes do not overlap.Use the **show conn detail** command to track how many DCD probes have been sent by the initiator and responder.

  For systems that are operating in a cluster or high-availability configuration, we recommend that you do not set the interval to less than one minute (0:1:0). If the connection needs to be moved between systems,

the changes required take longer than 30 seconds, and the connection might be deleted before the change is accomplished.

Step 5 To disable randomized sequence numbers, uncheck **Randomize Sequence Number**.

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.

Step 6 To customize TCP Normalizer behavior, check **Use TCP Map** and choose an existing TCP map from the drop-down list (if available), or add a new one by clicking **New**.

Step 7 To decrement time-to-live (TTL) on packets that match the class, check **Decrement time to live for a connection**.

Decrementing TTL is necessary for the ASA to show up in trace routes as one of the hops. You must also increase the rate limit for ICMP Unreachable messages on **Configuration > Device Management > Management Access > ICMP**.

Step 8 To enable TCP state bypass, check **TCP State Bypass**.

Step 9 To enable SCTP state bypass, check **SCTP State Bypass**.

Implement SCTP State Bypass to turn off SCTP stateful inspection. For more information, see SCTP Stateful Inspection, on page 375.

Step 10 (ASA on the Firepower 4100/9300 chassis, FXOS 1.1.3 or later, only.) To enable flow offload, check **Flow Offload**.

Eligible traffic is offloaded to a super fast path, where the flows are switched in the NIC itself. You must also enable the offload service. Select **Configuration** > **Firewall** > **Advanced** > **Offload Engine**.

Step 11 Click **OK** or **Finish**.

Configure TCP Options

You can configure options to control some aspects of TCP behavior. The defaults for these settings are appropriate for most networks.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Advanced** > **TCP Options**.

Step 2 Configure per-interface TCP reset behavior.

a) Select the interface you want to change and click **Edit**.

b) Select the desired options:

- **Send reset reply for denied inbound TCP packets**. Send TCP resets for all inbound TCP sessions that attempt to transit the ASA and are denied by the ASA based on access lists or AAA settings. The ASA also sends resets for packets that are allowed by an access list or AAA, but do not belong to an existing connection and are denied by the stateful firewall. Traffic between same security level interfaces is also affected. When this option is not enabled, the ASA silently discards denied packets.

c) Click **OK**.

Step 3 Configure other TCP options:

- **Send reset reply for denied outside TCP packets**. Send resets for TCP packets that terminate at the least secure interface and are denied by the ASA based on access lists or AAA settings. The ASA also sends resets for packets that are allowed by an access list or AAA, but do not belong to an existing connection and are denied by the stateful firewall. When this option is not enabled, the ASA silently discards the packets of denied connections.

  We recommend that you use this option with interface PAT. This allows the ASA to terminate the IDENT from an external SMTP or FTP server. Actively resetting these connections avoids the 30-second timeout delay.

- **Force maximum segment size for TCP connection to be X bytes**. Set the maximum TCP segment size in bytes, between 48 and any maximum number. The default value is 1380 bytes. You can disable this feature by setting bytes to 0.

- **Force minimum segment size for TCP connection to be X bytes**. Override the maximum segment size to be no less than bytes , between 48 and 65535 bytes. This feature is disabled by default (set to 0).

- **Force TCP connection to linger in TIME_WAIT state for at least 15 seconds after TCP close-down**. Forces each TCP connection to linger in a shortened TIME_WAIT state of at least 15 seconds after the final normal TCP close-down sequence. You might want to use this feature if an end host application default TCP terminating sequence is a simultaneous close.

- **TCP Maximum unprocessed segments**. Sets the maximum number of TCP unprocessed segments, from 6 to 24. The default is 6. If you find that SIP phones are not connecting to the call manager, you can try increasing the maximum number of unprocessed TCP segments.

Step 4 Click **Apply**.

Monitoring Connections

Use the following pages to monitor connections:

- **Home** > **Firewall Dashboard**, and look at the **Top Ten Protected Servers under SYN Attack** dashboard to monitor TCP Intercept. Click the **Detail** button to show history sampling data. The ASA samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

- **Monitoring** > **Properties** > **Connections**, to see current connections.

- **Monitoring** > **Properties** > **Connection Graphs**, to monitor performance.

In addition, you can enter the following commands using **Tools** > **Command Line Interface**.

- **show conn** [**detail**]

  Shows connection information. Detailed information uses flags to indicate special connection characteristics. For example, the “b” flag indicates traffic subject to TCP State Bypass.

  When you use the **detail** keyword, you can see information about Dead Connection Detection (DCD) probing, which shows how often the connection was probed by the initiator and responder. For example, the connection details for a DCD-enabled connection would look like the following:

```
TCP dmz: 10.5.4.11/5555 inside: 10.5.4.10/40299,
    flags UO , idle 1s, uptime 32m10s, timeout 1m0s, bytes 11828,
cluster sent/rcvd bytes 0/0, owners (0,255)
  Traffic received at interface dmz
        Locally received: 0 (0 byte/s)
  Traffic received at interface inside
        Locally received: 11828 (6 byte/s)
  Initiator: 10.5.4.10, Responder: 10.5.4.11
  DCD probes sent: Initiator 5, Responder 5
```

- **show flow-offload** {**info** [**detail**] | **cpu** | **flow** [**count** | **detail**] | **statistics**}

  Shows information about the flow offloading, including general status information, CPU usage for offloading, offloaded flow counts and details, and offloaded flow statistics.

- **show service-policy**

  Shows service policy statistics, including Dead Connection Detection (DCD) statistics.

- **show threat-detection statistics top tcp-intercept** [**all** | **detail**]

  View the top 10 protected servers under attack. The **all** keyword shows the history data of all the traced servers. The **detail** keyword shows history sampling data. The ASA samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

Note In the ASA configuration, embryonic connections—connection requests that have not yet completed the three-way handshake process—are closed quickly and not synchronized between the active and standby devices. This design ensures HA system efficiency and security. For this reason, there might be a difference in the number of connections on both ASAs, which is to be expected.

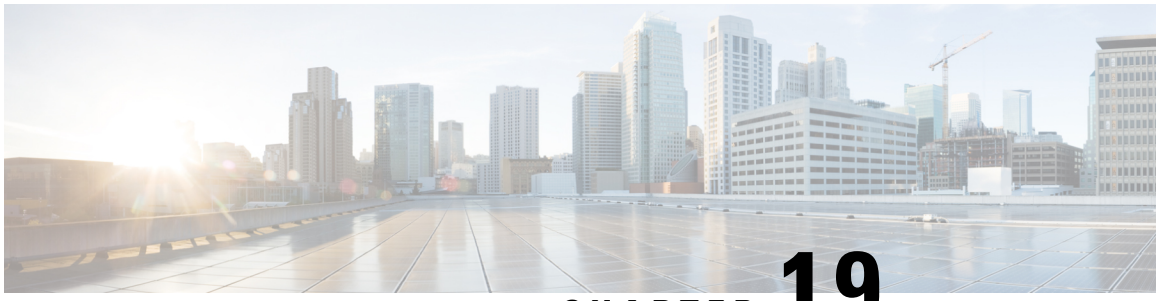
History for Connection Settings

Feature Name	Platform Releases	Description
TCP state bypass	8.2(1)	This feature was introduced. The following command was introduced: **set connection advanced-options tcp-state-bypass**.
Connection timeout for all protocols	8.2(2)	The idle timeout was changed to apply to all protocols, not just TCP. The following screen was modified: Configuration > Firewall > Service Policies > Rule Actions > Connection Settings.
Timeout for connections using a backup static route	8.2(5)/8.4(2)	When multiple static routes exist to a network with different metrics, the ASA uses the one with the best metric at the time of connection creation. If a better route becomes available, then this timeout lets connections be closed so a connection can be reestablished to use the better route. The default is 0 (the connection never times out). To take advantage of this feature, change the timeout to a new value. We modified the following screen: Configuration > Firewall > Advanced > Global Timeouts.

Feature Name	Platform Releases	Description
Configurable timeout for PAT xlate	8.4(3)	When a PAT xlate times out (by default after 30 seconds), and the ASA reuses the port for a new translation, some upstream routers might reject the new connection because the previous connection might still be open on the upstream device. The PAT xlate timeout is now configurable, to a value between 30 seconds and 5 minutes. We modified the following screen: Configuration > Firewall > Advanced > Global Timeouts. *This feature is not available in 8.5(1) or 8.6(1).*
Increased maximum connection limits for service policy rules	9.0(1)	The maximum number of connections for service policy rules was increased from 65535 to 2000000. We modified the following screen: Configuration > Firewall > Service Policy Rules > Connection Settings.
Decreased the half-closed timeout minimum value to 30 seconds	9.1(2)	The half-closed timeout minimum value for both the global timeout and connection timeout was lowered from 5 minutes to 30 seconds to provide better DoS protection. We modified the following screens: Configuration > Firewall > Service Policy Rules > Connection Settings; Configuration > Firewall > Advanced > Global Timeouts.
Connection holddown timeout for route convergence.	9.4(3) 9.6(2)	You can now configure how long the system should maintain a connection when the route used by the connection no longer exists or is inactive. If the route does not become active within this holddown period, the connection is freed. You can reduce the holddown timer to make route convergence happen more quickly. However, the 15 second default is appropriate for most networks to prevent route flapping. We modified the following screen: **Configuration** > **Firewall** > **Advanced** > **Global Timeouts**.
SCTP idle timeout and SCTP state bypass	9.5(2)	You can set an idle timeout for SCTP connections. You can also enable SCTP state bypass to turn off SCTP stateful inspection on a class of traffic. We modified the following screens: **Configuration** > **Firewall** > **Advanced** > **Global Timeouts**; **Configuration** > **Firewall** > **Service Policy Rules** wizard, **Connection Settings** tab.
Flow offload for the ASA on the Firepower 9300.	9.5(2.1)	You can identify flows that should be offloaded from the ASA and switched directly in the NIC (on the Firepower 9300). This provides improved performance for large data flows in data centers. This feature requires FXOS 1.1.3. We added or modified the following screens: **Configuration** > **Firewall** > **Advanced** > **Offload Engine**, the **Rule Actions** > **Connection Settings** tab when adding or editing rules under **Configuration** > **Firewall** > **Service Policy Rules**.

Feature Name	Platform Releases	Description
Flow offload support for the ASA on the Firepower 4100 series.	9.6(1)	You can identify flows that should be offloaded from the ASA and switched directly in the NIC for the Firepower 4100 series. This feature requires FXOS 1.1.4. There are no new commands or ASDM screens for this feature.
Flow offload support for multicast connections in transparent mode.	9.6(2)	You can now offload multicast connections to be switched directly in the NIC on transparent mode Firepower 4100 and 9300 series devices. Multicast offload is available for bridge groups that contain two and only two interfaces. There are no new commands or ASDM screens for this feature.
Changes in TCP option handling.	9.6(2)	You can now specify actions for the TCP MSS and MD5 options in a packet's TCP header when configuring a TCP map. In addition, the default handling of the MSS, timestamp, window-size, and selective-ack options has changed. Previously, these options were allowed, even if there were more than one option of a given type in the header. Now, packets are dropped by default if they contain more than one option of a given type. For example, previously a packet with 2 timestamp options would be allowed, now it will be dropped. You can configure a TCP map to allow multiple options of the same type for MD5, MSS, selective-ack, timestamp, and window-size. For the MD5 option, the previous default was to clear the option, whereas the default now is to allow it. You can also drop packets that contain the MD5 option. For the MSS option, you can set the maximum segment size in the TCP map (per traffic class). The default for all other TCP options remains the same: they are cleared. We modified the following screen: **Configuration** > **Firewall** > **Objects** > **TCP Maps** Add/Edit dialog box.
Stale route timeout for interior gateway protocols	9.7(1)	You can now configure the timeout for removing stale routes for interior gateway protocols such as OSPF. We modified the following screen: **Configuration** > **Firewall** > **Advanced** > **Global Timeouts**.
Global timeout for ICMP errors	9.8(1)	You can now set the idle time before the ASA removes an ICMP connection after receiving an ICMP echo-reply packet. When this timeout is disabled (the default), and you enable ICMP inspection, then the ASA removes the ICMP connection as soon as an echo-reply is received; thus any ICMP errors that are generated for the (now closed) connection are dropped. This timeout delays the removal of ICMP connections so you can receive important ICMP errors. We modified the following screen: **Configuration** > **Firewall** > **Advanced** > **Global Timeouts**.
Default idle timeout for TCP state bypass	9.10(1)	The default idle timeout for TCP state bypass connections is now 2 minutes instead of 1 hour.

Feature Name	Platform Releases	Description
Initiator and responder information for Dead Connection Detection (DCD), and DCD support in a cluster.	9.13(1)	If you enable Dead Connection Detection (DCD), you can use the **show conn detail** command to get information about the initiator and responder. Dead Connection Detection allows you to maintain an inactive connection, and the **show conn** output tells you how often the endpoints have been probed. In addition, DCD is now supported in a cluster. New/Modified screens: none.
Configure the maximum segment size (MSS) for embryonic connections.	9.16(1)	You can configure a service policy to set the server maximum segment size (MSS) for SYN-cookie generation for embryonic connections upon reaching the embryonic connections limit. This is meaningful for service policies where you are also setting embryonic connection maximums. New or changed screens: **Connection Settings** in the Add/Edit Service Policy wizard.
Configure the SIP Call Init idle global timeout	9.16(x)	You can configure the idle time after which a SIP session in Call_Init call state (after the initial invite message is received) will be deleted, i.e. when no response is received, between 0:0:1 and 1193:0:0. The default is 32 seconds (0:00:32).



CHAPTER 19

Quality of Service

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.

The following topics describe how to apply QoS policies.

- About QoS, on page 437
- Guidelines for QoS, on page 439
- Configure QoS, on page 439
- Monitor QoS, on page 443
- History for QoS, on page 445

About QoS

You should consider that in an ever-changing network environment, QoS is not a one-time deployment, but an ongoing, essential part of network design.

This section describes the QoS features available on the ASA.

Supported QoS Features

The ASA supports the following QoS features:

- Policing—To prevent classified traffic from hogging the network bandwidth, you can limit the maximum bandwidth used per class. See Policing, on page 438 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 Priority Queuing, on page 438.

What is a Token Bucket?

A token bucket is used to manage a device that regulates the data in a flow, for example, a traffic policer. 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 bytes per burst how much traffic can be sent within a given unit of time to not create scheduling concerns.
- 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 waits until the packet is discarded or marked down. 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.

Policing

Policing is a way of ensuring that no traffic exceeds the maximum rate (in bits/second) that you configure, thus ensuring that no one traffic class can take over the entire resource. When traffic exceeds the maximum rate, the ASA drops the excess traffic. Policing also sets the largest single burst of traffic allowed.

Priority Queuing

LLQ priority queuing lets you prioritize certain traffic flows (such as latency-sensitive traffic like voice and video) ahead of other traffic. Priority queuing uses an LLQ priority queue on an interface (see Configure the Priority Queue for an Interface, on page 441), while all other traffic goes into the “best effort” queue. Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped. This is called *tail drop*. To avoid having the queue fill up, you can increase the queue buffer size. You can also fine-tune the maximum number of packets allowed into the transmit queue. These options let you control the latency and robustness of the priority queuing. Packets in the LLQ queue are always transmitted before packets in the best effort queue.

How QoS Features Interact

You can configure each of the QoS features alone if desired for the ASA. Often, though, you configure multiple QoS features on the ASA so you can prioritize some traffic, for example, and prevent other traffic from causing bandwidth problems. You can configure:

Priority queuing (for specific traffic) + Policing (for the rest of the traffic).

You cannot configure priority queuing and policing for the same set of traffic.

DSCP (DiffServ) Preservation

DSCP (DiffServ) markings are preserved on all traffic passing through the ASA. The ASA does not locally mark/remark any classified traffic. For example, you could key off the Expedited Forwarding (EF) DSCP bits of every packet to determine if it requires “priority” handling and have the ASA direct those packets to the LLQ.

Guidelines for QoS

Context Mode Guidelines

Supported in single context mode only. Does not support multiple context mode.

Firewall Mode Guidelines

Supported in routed firewall mode only. Does not support transparent firewall mode.

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines and Limitations

- QoS is applied unidirectionally; only traffic that enters (or exits, depending on the QoS feature) the interface to which you apply the policy map is affected.
- For priority traffic, you cannot use the **class-default** class map.
- For priority queuing, the priority queue must be configured for a physical interface.
- For policing, to-the-box traffic is not supported.
- For policing, traffic to and from a VPN tunnel bypasses interface policing.
- For policing, when you match a tunnel group class map, only outbound policing is supported.

Configure QoS

Use the following sequence to implement QoS on the ASA.

Procedure

Step 1 Determine the Queue and TX Ring Limits for a Priority Queue, on page 440.
Step 2 Configure the Priority Queue for an Interface, on page 441.
Step 3 Configure a Service Rule for Priority Queuing and Policing, on page 442.

Determine the Queue and TX Ring Limits for a Priority Queue

Use the following worksheets to determine the priority queue and TX ring limits.

Queue Limit Worksheet

The following worksheet 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 Configure the Priority Queue for an Interface, on page 441.

Tips on the worksheet:

- Outbound bandwidth—For example, DSL might have an uplink speed of 768 Kbps. Check with your provider.
- Average packet size—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.
- Delay—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 15: Queue Limit Worksheet

1	_____ *Outbound bandwidth (Mbps or Kbps)*	Mbps	x	**125**	=	_____ *# of bytes/ms*		
		Kbps	x	**.125**	=	_____ *# of bytes/ms*		
2	_____ *# of bytes/ms from Step 1*		÷	_____ *Average packet size (bytes)*	x	_____ *Delay (ms)*	=	_____ *Queue limit (# of packets)*

TX Ring Limit Worksheet

The following worksheet 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.

Tips on the worksheet:

- Outbound bandwidth—For example, DSL might have an uplink speed of 768 Kbps. Check with your provider.
- Maximum packet size—Typically, the maximum size is 1538 bytes, or 1542 bytes for tagged Ethernet. If you allow jumbo frames (if supported for your platform), then the packet size might be larger.
- Delay—The delay depends on your application. For example, to control jitter for VoIP, you should use 20 ms.

Table 16: TX Ring Limit Worksheet

1	_____ *Outbound bandwidth (Mbps or Kbps)*	Mbps	x	**125**	=	_____ *# of bytes/ms*		
		Kbps	x	**0.125**	=	_____ *# of bytes/ms*		
2	_____ *# of bytes/ms from Step 1*		÷	_____ *Maximum packet size (bytes)*	x	_____ *Delay (ms)*	=	_____ *TX ring limit (# of packets)*

Configure the Priority Queue for an Interface

If you enable priority queuing for traffic on a physical interface, then you need to also create the priority queue on each interface. Each physical interface uses two queues: one for priority traffic, and the other for all other traffic. For the other traffic, you can optionally configure policing.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **Priority Queue**, and click **Add**.

Step 2 Configure the following options:

- **Interface**—The physical interface name on which you want to enable the priority queue, or for the ASASM, the VLAN interface name.
- **Queue Limit**—The number of average, 256-byte packets that the specified interface can transmit in a 500-ms interval. The range is 0-2048, and 2048 is the default.

  A packet that stays more than 500 ms in a network node might trigger a timeout in the end-to-end application. Such a packet can be discarded in each network node.

  Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped (called *tail drop*). To avoid having the queue fill up, you can use this option to increase the queue buffer size.

  The upper limit of the range of values for this option is determined dynamically at run time. The key determinants are the memory needed to support the queues and the memory available on the device.

The Queue Limit that you specify affects both the higher priority low-latency queue and the best effort queue.

- **Transmission Ring Limit**—The depth of the priority queues, which is the number of maximum 1550-byte packets that the specified interface can transmit in a 10-ms interval. The range is 3-511, and 511 is the default.

  This setting guarantees that the hardware-based transmit ring imposes no more than 10-ms of extra latency for a high-priority packet.

  This option sets the maximum number of low-latency or normal priority packets allowed into the Ethernet transmit driver before the driver pushes back to the queues on the interface to let them buffer packets until the congestion clears.

  The upper limit of the range of values is determined dynamically at run time. The key determinants are the memory needed to support the queues and the memory available on the device.

  The Transmission Ring Limit that you specify affects both the higher priority low-latency queue and the best-effort queue.

Step 3 Click **OK**, then **Apply**.

Configure a Service Rule for Priority Queuing and Policing

You can configure priority queuing and policing for different class maps within the same policy map. See How QoS Features Interact, on page 439 for information about valid QoS configurations.

Before you begin

- You cannot use the **class-default** class map for priority traffic.
- For policing, to-the-box traffic is not supported.
- For policing, traffic to and from a VPN tunnel bypasses interface policing.
- For policing, when you match a tunnel group class map, only outbound policing is supported.
- 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.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy**, and open a rule.

You can configure QoS as part of a new service policy rule, or you can edit an existing service policy.

Step 2 Proceed through the wizard to the Rules page, selecting the interface (or global) and traffic matching criteria along the way.

For policing traffic, you can choose to police all traffic that you are not prioritizing, or you can limit the traffic to certain types.

Tip
If you use an ACL for traffic matching, policing is applied in the direction specified in the ACL only. That is, traffic going from the source to the destination is policed, but not the reverse.

Step 3 In the Rule Actions dialog box, click the **QoS** tab.

Step 4 Select **Enable priority for this flow**.

If this service policy rule is for an individual interface, ASDM automatically creates the priority queue for the interface (Configuration > Device Management > Advanced > Priority Queue; for more information, see Configure the Priority Queue for an Interface, on page 441). If this rule is for the global policy, then you need to manually add the priority queue to one or more interfaces *before* you configure the service policy rule.

Step 5 Select **Enable policing**, then check the **Input policing** or **Output policing** (or both) check boxes to enable the specified type of traffic policing. For each type of traffic policing, configure the following options:

- **Committed Rate**—The rate limit for this traffic class, from 8000-2000000000 bits per second. For example, to limit traffic to 5Mbps, enter 5000000.
- **Conform Action**—The action to take when the traffic is below the policing rate and burst size. You can drop or transmit the traffic. The default is to transmit the traffic.
- **Exceed Action**—The action to take when traffic exceeds the policing rate and burst size. You can drop or transmit packets that exceed the policing rate and burst size. The default is to drop excess packets.
- **Burst Rate**—The maximum number of instantaneous bytes allowed in a sustained burst before throttling to the conforming rate value, between 1000 and 512000000 bytes. Calculate the burst size calculated as 1/32 of the conform-rate in bytes. For example, the burst size for a 5Mbps rate would be 156250. The default is 1500, but the system can recalculate any value you enter as needed.

Step 6 Click **Finish**, then **Apply**.

Monitor QoS

The following topics explain how to monitor QoS.

To monitor QoS in ASDM, you can enter commands at the Command Line Interface tool.

QoS Police Statistics

To view the QoS statistics for traffic policing, use 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
```

QoS Priority Statistics

To view statistics for service policies implementing the **priority** command, use 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
```

“Aggregate drop” denotes the aggregated drop in this interface; “aggregate transmit” denotes the aggregated number of transmitted packets in this interface.

QoS Priority Queue Statistics

To display the priority-queue statistics for an interface, use the **show priority-queue statistics** command. 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.

```
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#
```

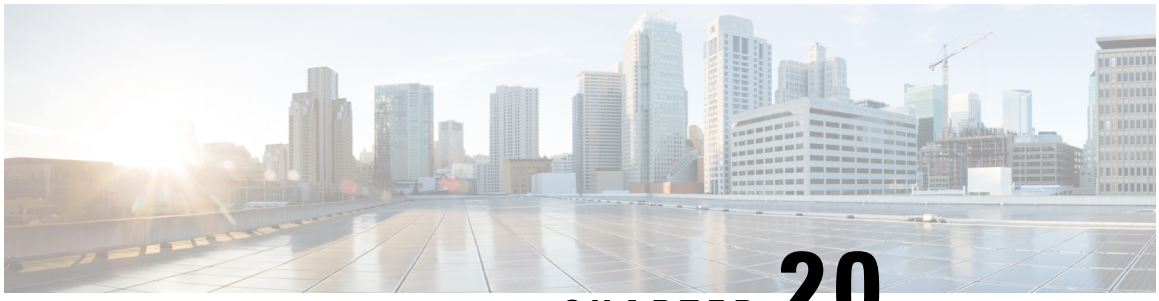
In this statistical report:

- “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.

History for QoS

Feature Name	Platform Releases	Description
Priority queuing and policing	7.0(1)	We introduced QoS priority queuing and policing. We introduced the following screens: Configuration > Device Management > Advanced > Priority Queue Configuration > Firewall > Service Policy Rules
Shaping and hierarchical priority queuing	7.2(4)/8.0(4)	We introduced QoS shaping and hierarchical priority queuing. We modified the following screen: Configuration > Firewall > Service Policy Rules.
Ten Gigabit Ethernet support for a standard priority queue on the ASA 5585-X	8.2(3)/8.4(1)	We added support for a standard priority queue on Ten Gigabit Ethernet interfaces for the ASA 5585-X.



CHAPTER 20

Threat Detection

The following topics describe how to configure threat detection statistics and scanning threat detection.

- Detecting Threats, on page 447
- Guidelines for Threat Detection, on page 449
- Defaults for Threat Detection, on page 450
- Configure Threat Detection, on page 451
- Monitoring Threat Detection, on page 453
- History for Threat Detection, on page 454

Detecting Threats

Threat detection on the ASA provides a front-line defense against attacks. Threat detection works at Layer 3 and 4 to develop a baseline for traffic on the device, analyzing packet drop statistics and accumulating "top" reports based on traffic patterns. In comparison, a module that provides IPS or Next Generation IPS services identifies and mitigates attack vectors up to Layer 7 on traffic the ASA permitted, and cannot see the traffic dropped already by the ASA. Thus, threat detection and IPS can work together to provide a more comprehensive threat defense.

Threat detection consists of the following elements:

- Different levels of statistics gathering for various threats.

  Threat detection statistics can help you manage threats to your ASA; for example, if you enable scanning threat detection, then viewing statistics can help you analyze the threat. You can configure two types of threat detection statistics:

  - Basic threat detection statistics—Includes information about attack activity for the system as a whole. Basic threat detection statistics are enabled by default and have no performance impact. See Basic Threat Detection Statistics, on page 448.
  - Advanced threat detection statistics—Tracks activity at an object level, so the ASA can report activity for individual hosts, ports, protocols, or ACLs. Advanced threat detection statistics can have a major performance impact, depending on the statistics gathered, so only the ACL statistics are enabled by default. See Advanced Threat Detection Statistics, on page 448.

- Scanning threat detection, which determines when a host is performing a scan. You can optionally shun any hosts determined to be a scanning threat. See Scanning Threat Detection, on page 449.

Basic Threat Detection Statistics

Using basic threat detection statistics, the ASA monitors the rate of dropped packets and security events due to the following reasons:

- Denial by ACLs.
- Bad packet format (such as invalid-ip-header or invalid-tcp-hdr-length).
- Connection limits exceeded (both system-wide resource limits, and limits set in the configuration).
- DoS attack detected (such as an invalid SPI, Stateful Firewall check failure).
- Basic firewall checks failed. This option is a combined rate that includes all firewall-related packet drops in this list. It does not include non-firewall-related drops such as interface overload, packets failed at application inspection, and scanning attack detected.
- Suspicious ICMP packets detected.
- Packets failed application inspection.
- Interface overload.
- Scanning attack detected. This option monitors scanning attacks; for example, the first TCP packet is not a SYN packet, or the TCP connection failed the 3-way handshake. Full scanning threat detection takes this scanning attack rate information and acts on it by classifying hosts as attackers and automatically shunning them, for example.
- Incomplete session detection such as TCP SYN attack detected or UDP session with no return data attack detected.

When the ASA detects a threat, it immediately sends a system log message (733100). The ASA tracks two types of rates: the average event rate over an interval, and the burst event rate over a shorter burst interval. The burst rate interval is 1/30th of the average rate interval or 10 seconds, whichever is higher. For each received event, the ASA checks the average and burst rate limits; if both rates are exceeded, then the ASA sends two separate system messages, with a maximum of one message for each rate type per burst period.

Basic threat detection affects performance only when there are drops or potential threats; even in this scenario, the performance impact is insignificant.

Advanced Threat Detection Statistics

Advanced threat detection statistics show both allowed and dropped traffic rates for individual objects such as hosts, ports, protocols, or ACLs.



Caution Enabling advanced statistics can affect the ASA performance, depending on the type of statistics enabled. Enabling host statistics affects performance in a significant way; if you have a high traffic load, you might consider enabling this type of statistics temporarily. Port statistics, however, has modest impact.

Scanning Threat Detection

A typical scanning attack consists of a host that tests the accessibility of every IP address in a subnet (by scanning through many hosts in the subnet or sweeping through many ports in a host or subnet). The scanning threat detection feature determines when a host is performing a scan. Unlike IPS scan detection that is based on traffic signatures, ASA threat detection scanning maintains an extensive database that contains host statistics that can be analyzed for scanning activity.

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

If the scanning threat rate is exceeded, then the ASA sends a syslog message (733101), and optionally shuns the attacker. The ASA tracks two types of rates: the average event rate over an interval, and the burst event rate over a shorter burst interval. The burst event rate is 1/30th of the average rate interval or 10 seconds, whichever is higher. For each event detected that is considered to be part of a scanning attack, the ASA checks the average and burst rate limits. If either rate is exceeded for traffic sent from a host, then that host is considered to be an attacker. If either rate is exceeded for traffic received by a host, then that host is considered to be a target.

The following table lists the default rate limits for scanning threat detection.

Table 17: Default Rate Limits for Scanning Threat Detection

Average Rate	Burst Rate
5 drops/sec over the last 600 seconds.	10 drops/sec over the last 20 second period.
5 drops/sec over the last 3600 seconds.	10 drops/sec over the last 120 second period.

Caution The scanning threat detection feature can affect the ASA performance and memory significantly while it creates and gathers host- and subnet-based data structure and information.

Guidelines for Threat Detection

Security Context Guidelines

Except for advanced threat statistics, threat detection is supported in single mode only. In Multiple mode, TCP Intercept statistics are the only statistic supported.

Types of Traffic Monitored

- For statistics, only through-the-box traffic is monitored; to-the-box traffic is not monitored.
- Traffic that is denied by an ACL does not trigger scanning threat detection; only traffic that is allowed through the ASA and that creates a flow is affected by scanning threat detection.

Defaults for Threat Detection

Basic threat detection statistics are enabled by default.

The following table lists the default settings. You can view all these default settings using the **show running-config all threat-detection** command in Tools > Command Line Interface.

For advanced statistics, by default, statistics for ACLs are enabled.

Table 18: Basic Threat Detection Default Settings

	Trigger Settings	
Packet Drop Reason	**Average Rate**	**Burst Rate**
• DoS attack detected • Bad packet format • Connection limits exceeded • Suspicious ICMP packets detected	100 drops/sec over the last 600 seconds.	400 drops/sec over the last 20 second period.
	80 drops/sec over the last 3600 seconds.	320 drops/sec over the last 120 second period.
Scanning attack detected	5 drops/sec over the last 600 seconds.	10 drops/sec over the last 20 second period.
	4 drops/sec over the last 3600 seconds.	8 drops/sec over the last 120 second period.
Incomplete session detected such as TCP SYN attack detected or UDP session with no return data attack detected (combined)	100 drops/sec over the last 600 seconds.	200 drops/sec over the last 20 second period.
	80 drops/sec over the last 3600 seconds.	160 drops/sec over the last 120 second period.
Denial by ACLs	400 drops/sec over the last 600 seconds.	800 drops/sec over the last 20 second period.
	320 drops/sec over the last 3600 seconds.	640 drops/sec over the last 120 second period.
• Basic firewall checks failed • Packets failed application inspection	400 drops/sec over the last 600 seconds.	1600 drops/sec over the last 20 second period.
	320 drops/sec over the last 3600 seconds.	1280 drops/sec over the last 120 second period.
Interface overload	2000 drops/sec over the last 600 seconds.	8000 drops/sec over the last 20 second period.
	1600 drops/sec over the last 3600 seconds.	6400 drops/sec over the last 120 second period.

Configure Threat Detection

Basic threat detection statistics are enabled by default, and might be the only threat detection service that you need. Use the following procedure if you want to implement additional threat detection services.

Procedure

Step 1 Configure Basic Threat Detection Statistics, on page 451.

Basic threat detection statistics include activity that might be related to an attack, such as a DoS attack.

Step 2 Configure Advanced Threat Detection Statistics, on page 451.

Step 3 Configure Scanning Threat Detection, on page 452.

Configure Basic Threat Detection Statistics

Basic threat detection statistics is enabled by default. You can disabled it, or turn it on again if you disable it.

Procedure

Step 1 Choose the **Configuration** > **Firewall** > **Threat Detection**.

Step 2 Select or deselect **Enable Basic Threat Detection** as desired.

Step 3 Click **Apply**.

Configure Advanced Threat Detection Statistics

You can configure the ASA to collect extensive statistics. By default, statistics for ACLs are enabled. To enable other statistics, perform the following steps.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Threat Detection**.

Step 2 In the Scanning Threat Statistics area, choose one of the following options:

- **Enable All Statistics**.
- **Disable All Statistics**.
- **Enable Only Following Statistics**.

Step 3 If you chose **Enable Only Following Statistics**, then select one or more of the following options:

- **Hosts**—Enables host statistics. The host statistics accumulate for as long as the host is active and in the scanning threat host database. The host is deleted from the database (and the statistics cleared) after 10 minutes of inactivity.
- **Access Rules** (enabled by default)—Enables statistics for access rules.
- **Port**—Enables statistics for TCP and UDP ports.
- **Protocol**—Enables statistics for non-TCP/UDP IP protocols.
- **TCP-Intercept**—Enables statistics for attacks intercepted by TCP Intercept (to enable TCP Intercept, see Protect Servers from a SYN Flood DoS Attack (TCP Intercept), on page 417).

Step 4 For host, port, and protocol statistics, you can change the number of rate intervals collected. In the Rate Intervals area, choose **1 hour**, **1 and 8 hours**, or **1, 8 and 24 hours** for each statistics type. The default interval is **1 hour**, which keeps the memory usage low.

Step 5 For TCP Intercept statistics, you can set the following options in the TCP Intercept Threat Detection area:

- **Monitoring Window Size**—Sets the size of the history monitoring window, between 1 and 1440 minutes. The default is 30 minutes. The ASA samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.
- **Burst Threshold Rate**—Sets the threshold for syslog message generation, between 25 and 2147483647. The default is 400 per second. When the burst rate is exceeded, syslog message 733104 is generated.
- **Average Threshold Rate**—Sets the average rate threshold for syslog message generation, between 25 and 2147483647. The default is 200 per second. When the average rate is exceeded, syslog message 733105 is generated.

Click **Set Default** to restore the default values.

Step 6 Click **Apply**.

Configure Scanning Threat Detection

You can configure scanning threat detection to identify attackers and optionally shun them.

You can configure the ASA to send system log messages about an attacker or you can automatically shun the host. By default, the system log message 730101 is generated when a host is identified as an attacker. Be sure to except addresses from shunning when you expect a lot of messages from the host. For example, if you have enabled PIM multicast, exempt the PIM routers or PIM messages will be dropped.

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Threat Detection**.

Step 2 Select **Enable Scanning Threat Detection**.

Step 3 (Optional) To automatically terminate a host connection when the ASA identifies the host as an attacker, select **Shun Hosts detected by scanning threat** and fill in these options if desired:

- To exempt host IP addresses from being shunned, enter an address or the name of a network object in the **Networks excluded from shun** field. You can enter multiple addresses or subnets separated by commas. To choose a network from the list of IP address objects, click the **...** button.
- To set the duration of a shun for an attacking host, select **Set Shun Duration** and enter a value between 10 and 2592000 seconds. The default length is 3600 seconds (1 hour). To restore the default value, click **Set Default**.

Step 4 Click **Apply**.

Monitoring Threat Detection

The following topics explain how to monitor threat detection and view traffic statistics.

Monitoring Basic Threat Detection Statistics

Choose **Home > Firewall Dashboard > Traffic Overview** to view basic threat detection statistics.

Monitoring Advanced Threat Detection Statistics

You can monitor advanced threat statistics using the following dashboards:

- **Home > Firewall Dashboard > Top 10 Access Rules**—Displays the most hit access rules. Permits and denies are not differentiated in this graph. You can track denied traffic in the **Traffic Overview > Dropped Packets Rate** graph.
- **Home > Firewall Dashboard > Top Usage Statistics**—The **Top 10 Sources** and **Top 10 Destinations** tabs show statistics for hosts. Due to the threat detection algorithm, an interface used as a combination failover and state link could appear in the top 10 hosts; this is expected behavior, and you can ignore this IP address in the display.

  The **Top 10 Services** tab shows statistics for both ports and protocols (both must be enabled for the display), and shows the combined statistics of TCP/UDP port and IP protocol types. TCP (protocol 6) and UDP (protocol 17) are not included in the display for IP protocols; TCP and UDP ports are, however, included in the display for ports. If you only enable statistics for one of these types, port or protocol, then you will only view the enabled statistics.
- **Home > Firewall Dashboard > Top Ten Protected Servers under SYN Attack**—Shows the TCP Intercept statistics. Click the **Detail** button to show history sampling data. The ASA samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

History for Threat Detection

Feature Name	Platform Releases	Description
Basic and advanced threat detection statistics, scanning threat detection	8.0(2)	Basic and advanced threat detection statistics, scanning threat detection was introduced. The following screens were introduced: **Configuration > Firewall > Threat Detection**, **Home > Firewall Dashboard > Traffic Overview**, **Home > Firewall Dashboard > Top 10 Access Rules**, **Home > Firewall Dashboard > Top Usage Status**, **Home > Firewall Dashboard > Top 10 Protected Servers Under SYN Attack**.
Shun duration	8.0(4)/8.1(2)	You can now set the shun duration, The following screens was modified: **Configuration > Firewall > Threat Detection**.
TCP Intercept statistics	8.0(4)/8.1(2)	TCP Intercept statistics were introduced. The following screens were introduced or modified: **Configuration > Firewall > Threat Detection**, **Home > Firewall Dashboard > Top 10 Protected Servers Under SYN Attack**.
Customize host statistics rate intervals	8.1(2)	You can now customize the number of rate intervals for which statistics are collected. The default number of rates was changed from 3 to 1. The following screen was modified: **Configuration > Firewall > Threat Detection**.
Burst rate interval changed to 1/30th of the average rate.	8.2(1)	In earlier releases, the burst rate interval was 1/60th of the average rate. To maximize memory usage, the sampling interval was reduced to 30 times during the average rate.
Customize port and protocol statistics rate intervals	8.3(1)	You can now customize the number of rate intervals for which statistics are collected. The default number of rates was changed from 3 to 1. The following screen was modified: **Configuration > Firewall > Threat Detection**.
Improved memory usage	8.3(1)	The memory usage for threat detection was improved.