

Cisco ASA with FirePOWER Services Local Management Configuration Guide

Version 6.2.2
August 16, 2017

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

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

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

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

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

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

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

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.

© 2017 Cisco Systems, Inc. All rights reserved.



CONTENTS

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

Introduction to the ASA FirePOWER Module 1-1

ASA FirePOWER Module Components 1-2

Access Control 1-2

Intrusion Detection and Prevention 1-2

Advanced Malware Protection and File Control 1-3

Application Programming Interfaces 1-3

License Conventions 1-3

IP Address Conventions 1-4

CHAPTER 2 **Managing Reusable Objects** 2-1

Using the Object Manager 2-2

Grouping Objects 2-2

Browsing, Sorting, and Filtering Objects 2-3

Working with Network Objects 2-3

Working with Security Intelligence Lists and Feeds 2-4

Working with the Global Whitelist and Blacklist 2-6

Working with the Intelligence Feed 2-6

Working with Custom Security Intelligence Feeds 2-7

Manually Updating Security Intelligence Feeds 2-7

Working with Custom Security Intelligence Lists 2-8

Working with Port Objects 2-9

Working with URL Objects 2-10

Working with Application Filters 2-11

Working with Variable Sets 2-13

Optimizing Predefined Default Variables 2-13

Understanding Variable Sets 2-16

Managing Variable Sets 2-17

Managing Variables 2-18

Adding and Editing Variables 2-20

Resetting Variables 2-25

Linking Variable Sets to Intrusion Policies 2-26

Understanding Advanced Variables 2-27
Working with Sinkhole Objects 2-27
Working with File Lists 2-28
Uploading Multiple SHA-256 Values to a File List 2-28
Uploading an Individual File to a File List 2-30
Adding a SHA-256 Value to the File List 2-30
Modifying Files on a File List 2-31
Downloading a Source File from a File List 2-31
Working with Security Zones 2-32
Working with Cipher Suite Lists 2-32
Working with Distinguished Name Objects 2-33
Working with PKI Objects 2-35
Working with Internal Certificate Authority Objects 2-35
Working with Trusted Certificate Authority Objects 2-39
Working with External Certificate Objects 2-41
Working with Internal Certificate Objects 2-41
Working with Geolocation Objects 2-42
Working with Security Group Tag Objects 2-43

CHAPTER 3 **Managing Device Configuration** 3-1
Editing Device Configuration 3-1
Editing General Device Configuration 3-1
Viewing Device System Settings 3-2
Understanding Advanced Device Settings 3-2
Editing Advanced Device Settings 3-3
Managing ASA FirePOWER Module Interfaces 3-4
Applying Changes to Device Configuration 3-4
Using the Device Management Revision Comparison Report 3-5
Configuring Remote Management 3-5
Editing Remote Management 3-7
Configuring eStreamer on the eStreamer Server 3-7

CHAPTER 4 **Getting Started with Access Control Policies** 4-1
Access Control License and Role Requirements 4-2
License Requirements for Access Control 4-2
Creating a Basic Access Control Policy 4-3
Setting Default Handling and Inspection for Network Traffic 4-4
Managing Access Control Policies 4-6

Editing Access Control Policies 4-7
Associating Other Policies with Access Control 4-9
Understanding Out-of-Date Policy Warnings 4-10
Deploying Configuration Changes 4-11
Troubleshooting Access Control Policies and Rules 4-12
Simplifying Rules to Improve Performance 4-13
Understanding Rule Preemption and Invalid Configuration Warnings 4-14
Ordering Rules to Improve Performance and Avoid Preemption 4-14
Generating a Report of Current Access Control Settings 4-15
Comparing Access Control Policies 4-16

CHAPTER 5 **Blacklisting Using Security Intelligence IP Address Reputation 5-1**
Choosing a Security Intelligence Strategy 5-2
Building the Security Intelligence Whitelist and Blacklist 5-3
Searching for Objects to Whitelist or Blacklist 5-5

CHAPTER 6 **Tuning Traffic Flow Using Access Control Rules 6-1**
Creating and Editing Access Control Rules 6-2
Specifying a Rule's Order of Evaluation 6-4
Using Conditions to Specify the Traffic a Rule Handles 6-5
Using Rule Actions to Determine Traffic Handling and Inspection 6-7
Adding Comments to a Rule 6-10
Managing Access Control Rules in a Policy 6-11
Searching Access Control Rules 6-12
Enabling and Disabling Rules 6-13
Changing a Rule's Position or Category 6-13

CHAPTER 7 **Controlling Traffic with Network-Based Rules 7-1**
Controlling Traffic by Security Zone 7-1
Controlling Traffic by Network or Geographical Location 7-3
Controlling Traffic by Port and ICMP Codes 7-5

CHAPTER 8 **Controlling Traffic with Reputation-Based Rules 8-1**
Controlling Application Traffic 8-2
Matching Traffic with Application Filters 8-3
Matching Traffic from Individual Applications 8-4
Adding an Application Condition to an Access Control Rule 8-5
Limitations to Application Control 8-6

Blocking URLs 8-7
Performing Reputation-Based URL Blocking 8-8
Performing Manual URL Blocking 8-10
Limitations to URL Detection and Blocking 8-11
Allowing Users to Bypass URL Blocks 8-12
Displaying a Custom Web Page for Blocked URLs 8-14

CHAPTER 9 **Access Control Rules: Realms and Users 9-1**
Realm, User, User Group, and ISE Attribute Access Control Rule Conditions 9-1
Troubleshooting Issues with User Access Control Rules 9-2
Adding a Realm, User, or User Group Condition to an Access Control Rule 9-3
Configuring ISE Attribute Conditions 9-3

CHAPTER 10 **Access Control Rules: Custom Security Group Tags 10-1**
ISE SGT v. Custom SGT Rule Conditions 10-1
Automatic Transition from Custom SGT to ISE SGT Rule Conditions 10-2
Configuring Custom SGT Conditions 10-2
Troubleshooting Custom SGT Conditions 10-3
10-3

CHAPTER 11 **Controlling Traffic Using Intrusion and File Policies 11-1**
Inspecting Allowed Traffic For Intrusions and Malware 11-2
Understanding File and Intrusion Inspection Order 11-2
Configuring an Access Control Rule to Perform AMP or File Control 11-3
Configuring an Access Control Rule to Perform Intrusion Prevention 11-4
Tuning Intrusion Prevention Performance 11-6
Limiting Pattern Matching for Intrusions 11-6
Overriding Regular Expression Limits for Intrusion Rules 11-7
Limiting Intrusion Events Generated Per Packet 11-8
Configuring Packet and Intrusion Rule Latency Thresholds 11-9
Configuring Intrusion Performance Statistic Logging 11-15
Tuning File and Malware Inspection Performance and Storage 11-16

CHAPTER 12 **Intelligent Application Bypass (IAB) 12-1**
Introduction to IAB 12-1
IAB Options 12-2
Configuring IAB 12-4

IAB Logging and Analysis 12-5

CHAPTER 13 **Access Control Using Content Restriction 13-1**

Using Access Control Rules to Enforce Content Restriction 13-1

Safe Search Options for Access Control Rules 13-3

YouTube EDU Options for Access Control Rules 13-3

Content Restriction Rule Order 13-4

13-4

CHAPTER 14 **Understanding Traffic Decryption 14-1**

SSL Handshake Processing 14-2

ClientHello Message Handling 14-2

ServerHello and Server Certificate Message Handling 14-4

SSL Inspection Requirements 14-5

Deploying ASA FirePOWER Modules that Support SSL Inspection 14-6

License Requirements for SSL Inspection 14-6

Collecting Prerequisite Information to Configure SSL Rules 14-7

Analyzing SSL Inspection Appliance Deployments 14-7

Example: Decrypting Traffic in a Passive Deployment 14-8

Example: Decrypting Traffic in an Inline Deployment 14-11

CHAPTER 15 **Getting Started with SSL Policies 15-1**

Creating a Basic SSL Policy 15-2

Setting Default Handling and Inspection for Encrypted Traffic 15-3

Setting Default Handling for Undecryptable Traffic 15-4

Editing an SSL Policy 15-6

Applying Decryption Settings Using Access Control 15-8

Generating a Report of Current Traffic Decryption Settings 15-9

Comparing SSL Policies 15-10

CHAPTER 16 **Getting Started with SSL Rules 16-1**

Configuring Supporting Inspection Information 16-3

Understanding and Creating SSL Rules 16-4

Specifying an SSL Rule's Order of Evaluation 16-6

Using Conditions to Specify the Encrypted Traffic a Rule Handles 16-6

Using Rule Actions to Determine Encrypted Traffic Handling and Inspection 16-8

Monitor Action: Postponing Action and Ensuring Logging 16-9

Do Not Decrypt Action: Passing Encrypted Traffic Without Inspection 16-9
Blocking Actions: Blocking Encrypted Traffic Without Inspection 16-9
Decrypt Actions: Decrypting Traffic for Further Inspection 16-10
Managing SSL Rules in a Policy 16-12
Searching SSL Rules 16-12
Enabling and Disabling SSL Rules 16-13
Changing an SSL Rule's Position or Category 16-13
Troubleshooting SSL Rules 16-15
Configuring SSL Inspection to Improve Performance 16-18

CHAPTER 17 **Tuning Traffic Decryption Using SSL Rules 17-1**
Controlling Encrypted Traffic with Network-Based Conditions 17-1
Controlling Encrypted Traffic by Network Zone 17-2
Controlling Encrypted Traffic by Network or Geographical Location 17-3
Controlling Encrypted Traffic by Port 17-5
Controlling Encrypted Traffic Based on User 17-6
Controlling Encrypted Traffic by Reputation 17-7
Controlling Encrypted Traffic Based on Application 17-8
Controlling Encrypted Traffic by URL Category and Reputation 17-13
Controlling Traffic Based on Server Certificate Characteristics 17-16
Controlling Encrypted Traffic by Certificate Distinguished Name 17-17
Controlling Encrypted Traffic by Certificate 17-19
Controlling Encrypted Traffic by Certificate Status 17-20
Controlling Encrypted Traffic by Cipher Suite 17-25
Controlling Traffic by Encryption Protocol Version 17-26

CHAPTER 18 **Understanding Network Analysis and Intrusion Policies 18-1**
Understanding How Policies Examine Traffic For Intrusions 18-2
Decoding, Normalizing, and Preprocessing: Network Analysis Policies 18-3
Access Control Rules: Intrusion Policy Selection 18-4
Intrusion Inspection: Intrusion Policies, Rules, and Variable Sets 18-5
Intrusion Event Generation 18-6
Comparing System-Provided with Custom Policies 18-7
Understanding the System-Provided Policies 18-8
Benefits of Custom Policies 18-9
Benefits of a Custom Network Analysis Policy 18-9
Benefits of Custom Intrusion Policies 18-10
Limitations of Custom Policies 18-11
Using the Navigation Panel 18-13

Resolving Conflicts and Committing Policy Changes 18-15

CHAPTER 19 **Using Layers in a Network Analysis or Intrusion Policy 19-1**

Understanding the Layer Stack 19-1

Understanding the Base Layer 19-2

Managing Layers 19-5

Adding a Layer 19-7

Changing a Layer's Name and Description 19-7

Moving, Copying, and Deleting Layers 19-8

Merging Layers 19-8

Sharing Layers Between Policies 19-9

Configuring Intrusion Rules in Layers 19-11

Configuring Preprocessors and Advanced Settings in Layers 19-14

CHAPTER 20 **Customizing Traffic Preprocessing 20-1**

Setting the Default Intrusion Policy for Access Control 20-1

Customizing Preprocessing with Network Analysis Policies 20-2

Setting the Default Network Analysis Policy for Access Control 20-3

Specifying Traffic to Preprocess Using Network Analysis Rules 20-4

Managing Network Analysis Rules 20-7

CHAPTER 21 **Getting Started with Network Analysis Policies 21-1**

Creating a Custom Network Analysis Policy 21-2

Managing Network Analysis Policies 21-3

Editing Network Analysis Policies 21-3

Allowing Preprocessors to Affect Traffic in Inline Deployments 21-5

Configuring Preprocessors in a Network Analysis Policy 21-6

Generating a Report of Current Network Analysis Settings 21-8

Comparing Two Network Analysis Policies or Revisions 21-9

CHAPTER 22 **Using Application Layer Preprocessors 22-1**

CHAPTER 22

Decoding DCE/RPC Traffic 22-2

Selecting Global DCE/RPC Options 22-3

Understanding Target-Based DCE/RPC Server Policies 22-4

Understanding DCE/RPC Transports 22-5

Selecting DCE/RPC Target-Based Policy Options 22-8

Configuring the DCE/RPC Preprocessor 22-11

Detecting Exploits in DNS Name Server Responses 22-14
Understanding DNS Preprocessor Resource Record Inspection 22-14
Detecting Overflow Attempts in RData Text Fields 22-15
Detecting Obsolete DNS Resource Record Types 22-16
Detecting Experimental DNS Resource Record Types 22-16
Configuring the DNS Preprocessor 22-17
Decoding FTP and Telnet Traffic 22-18
Understanding Global FTP and Telnet Options 22-18
Configuring Global FTP/Telnet Options 22-19
Understanding Telnet Options 22-20
Configuring Telnet Options 22-21
Understanding Server-Level FTP Options 22-22
Configuring Server-Level FTP Options 22-25
Understanding Client-Level FTP Options 22-28
Configuring Client-Level FTP Options 22-29
Decoding HTTP Traffic 22-31
Selecting Global HTTP Normalization Options 22-31
Configuring Global HTTP Configuration Options 22-32
Selecting Server-Level HTTP Normalization Options 22-33
Selecting Server-Level HTTP Normalization Encoding Options 22-41
Configuring HTTP Server Options 22-43
Enabling Additional HTTP Inspect Preprocessor Rules 22-45
Using the Sun RPC Preprocessor 22-46
Configuring the Sun RPC Preprocessor 22-47
Decoding the Session Initiation Protocol 22-48
Selecting SIP Preprocessor Options 22-49
Configuring the SIP Preprocessor 22-50
Enabling Additional SIP Preprocessor Rules 22-51
Configuring the GTP Command Channel 22-52
Decoding IMAP Traffic 22-54
Selecting IMAP Preprocessor Options 22-54
Configuring the IMAP Preprocessor 22-55
Enabling Additional IMAP Preprocessor Rules 22-56
Decoding POP Traffic 22-57
Selecting POP Preprocessor Options 22-57
Configuring the POP Preprocessor 22-58
Enabling Additional POP Preprocessor Rules 22-59
Decoding SMTP Traffic 22-60
Understanding SMTP Decoding 22-60

Configuring SMTP Decoding 22-64
Enabling SMTP Maximum Decoding Memory Alerting 22-67
Detecting Exploits Using the SSH Preprocessor 22-67
Selecting SSH Preprocessor Options 22-68
Configuring the SSH Preprocessor 22-70
Using the SSL Preprocessor 22-71
Understanding SSL Preprocessing 22-71
Enabling SSL Preprocessor Rules 22-72
Configuring the SSL Preprocessor 22-73

CHAPTER 23 **Configuring SCADA Preprocessing** 23-1
Configuring the Modbus Preprocessor 23-1
Configuring the DNP3 Preprocessor 23-3

CHAPTER 24 **Configuring Transport & Network Layer Preprocessing** 24-1
Configuring Advanced Transport/Network Settings 24-1
Initiating Active Responses with Intrusion Drop Rules 24-2
Troubleshooting: Logging Session Termination Messages 24-3
Verifying Checksums 24-4
Normalizing Inline Traffic 24-6
Defragmenting IP Packets 24-11
Understanding IP Fragmentation Exploits 24-11
Target-Based Defragmentation Policies 24-12
Selecting Defragmentation Options 24-13
Configuring IP Defragmentation 24-14
Understanding Packet Decoding 24-16
Configuring Packet Decoding 24-19
Using TCP Stream Preprocessing 24-20
Understanding State-Related TCP Exploits 24-20
Selecting The TCP Global Option 24-21
Understanding Target-Based TCP Policies 24-21
Selecting TCP Policy Options 24-22
Reassembling TCP Streams 24-26
Configuring TCP Stream Preprocessing 24-28
Using UDP Stream Preprocessing 24-31
Configuring UDP Stream Preprocessing 24-31

CHAPTER 25 **Tuning Preprocessing in Passive Deployments** 25-1
Understanding Adaptive Profiles 25-1
Using Adaptive Profiles with Preprocessors 25-1
Configuring Adaptive Profiles 25-2

CHAPTER 26 **Getting Started with Intrusion Policies** 26-1
Creating a Custom Intrusion Policy 26-2
Managing Intrusion Policies 26-3
Editing Intrusion Policies 26-4
Setting Drop Behavior in an Inline Deployment 26-5
Configuring Advanced Settings in an Intrusion Policy 26-6
Applying an Intrusion Policy 26-7
Generating a Report of Current Intrusion Settings 26-8
Comparing Two Intrusion Policies or Revisions 26-9

CHAPTER 27 **Tuning Intrusion Policies Using Rules** 27-1
Understanding Intrusion Prevention Rule Types 27-1
Viewing Rules in an Intrusion Policy 27-2
Sorting the Rule Display 27-4
Viewing Rule Details 27-4
Filtering Rules in an Intrusion Policy 27-9
Understanding Rule Filtering in an Intrusion Policy 27-9
Setting a Rule Filter in an Intrusion Policy 27-17
Setting Rule States 27-19
Filtering Intrusion Event Notification Per Policy 27-21
Configuring Event Thresholding 27-21
Configuring Suppression Per Intrusion Policy 27-25
Adding Dynamic Rule States 27-28
Understanding Dynamic Rule States 27-28
Setting a Dynamic Rule State 27-29
Adding SNMP Alerts 27-31
Adding Rule Comments 27-32

CHAPTER 28 **Detecting Specific Threats** 28-1
Detecting Back Orifice 28-1
Detecting Portscans 28-3
Configuring Portscan Detection 28-5

Understanding Portscan Events 28-7
Preventing Rate-Based Attacks 28-9
Understanding Rate-Based Attack Prevention 28-9
Rate-Based Attack Prevention and Other Filters 28-12
Configuring Rate-Based Attack Prevention 28-17
Detecting Sensitive Data 28-19
Deploying Sensitive Data Detection 28-20
Selecting Global Sensitive Data Detection Options 28-20
Selecting Individual Data Type Options 28-21
Using Predefined Data Types 28-22
Configuring Sensitive Data Detection 28-23
Selecting Application Protocols to Monitor 28-25
Special Case: Detecting Sensitive Data in FTP Traffic 28-26
Using Custom Data Types 28-27

CHAPTER 29 **Globally Limiting Intrusion Event Logging** 29-1
Understanding Thresholding 29-1
Understanding Thresholding Options 29-2
Configuring Global Thresholds 29-3
Disabling the Global Threshold 29-4

CHAPTER 30 **Understanding and Writing Intrusion Rules** 30-1
Understanding Rule Anatomy 30-2
Understanding Rule Headers 30-3
Specifying Rule Actions 30-4
Specifying Protocols 30-4
Specifying IP Addresses In Intrusion Rules 30-5
Defining Ports in Intrusion Rules 30-8
Specifying Direction 30-9
Understanding Keywords and Arguments in Rules 30-9
Defining Intrusion Event Details 30-11
Searching for Content Matches 30-15
Constraining Content Matches 30-17
Replacing Content in Inline Deployments 30-29
Using Byte_Jump and Byte_Test 30-30
Searching for Content Using PCRE 30-35
Adding Metadata to a Rule 30-43
Inspecting IP Header Values 30-45
Inspecting ICMP Header Values 30-48

Inspecting TCP Header Values and Stream Size 30-50
Enabling and Disabling TCP Stream Reassembly 30-54
Extracting SSL Information from a Session 30-54
Inspecting Application Layer Protocol Values 30-56
Inspecting Packet Characteristics 30-79
Reading Packet Data into Keyword Arguments 30-82
Initiating Active Responses with Rule Keywords 30-84
Filtering Events 30-87
Evaluating Post-Attack Traffic 30-88
Detecting Attacks That Span Multiple Packets 30-89
Generating Events on the HTTP Encoding Type and Location 30-94
Detecting File Types and Versions 30-96
Pointing to a Specific Payload Type 30-98
Pointing to the Beginning of the Packet Payload 30-99
Decoding and Inspecting Base64 Data 30-99
Constructing a Rule 30-101
Writing New Rules 30-101
Modifying Existing Rules 30-103
Adding Comments to Rules 30-104
Deleting Custom Rules 30-105
Filtering Rules on the Rule Editor Page 30-105
Using Keywords in a Rule Filter 30-106
Using Character Strings in a Rule Filter 30-107
Combining Keywords and Character Strings in a Rule Filter 30-108
Filtering Rules 30-108

CHAPTER 31 **Introduction to Identity Data** 31-1
Uses for Identity Data 31-1
User Detection Fundamentals 31-1
The User Activity Database 31-3
The Users Database 31-3
Current User Identities 31-3
User Database Limits 31-3

CHAPTER 32 **Realms and Identity Policies** 32-1
Realm Fundamentals 32-1
Supported Servers for Realms 32-2
Supported Server Field Names 32-3
Troubleshooting Issues with Realms 32-3

Identity Policy Fundamentals 32-4
Creating a Realm 32-4
Realm Fields 32-5
Configuring Basic Realm Information 32-7
Configuring a Realm Directory 32-7
Configuring an Identity Policy 32-9
Managing Realms 32-17
Managing the Identity Policy 32-19

CHAPTER 33 **User Identity Sources** 33-1
Troubleshooting Issues with User Identity Sources 33-1
The User Agent Identity Source 33-3
Configuring a User Agent Connection 33-3
The ISE/ISE-PIC Identity Source 33-4
ISE/ISE-PIC Fields 33-5
Configuring an ISE/ISE-PIC Connection 33-6
The Captive Portal Active Authentication Identity Source 33-7
ASA FirePOWER Module-Server Downloads 33-7

CHAPTER 34 **DNS Policies** 34-1
DNS Policy Overview 34-1
DNS Policy Components 34-1
Editing a DNS Policy 34-2
DNS Rules 34-2
Creating and Editing DNS Rules 34-3
DNS Rule Management 34-4
DNS Policy Deploy 34-8

CHAPTER 35 **Blocking Malware and Prohibited Files** 35-1
Understanding Malware Protection and File Control 35-1
Configuring Malware Protection and File Control 35-3
Logging Events Based on Malware Protection and File Control 35-3
Understanding and Creating File Policies 35-4
Creating a File Policy 35-9
Working with File Rules 35-10
Configuring Advanced File Policy General Options 35-11
Comparing Two File Policies 35-12

CHAPTER 36 **Logging Connections in Network Traffic** 36-1
Deciding Which Connections To Log 36-1
Logging Critical Connections 36-2
Logging the Beginning and End of Connections 36-3
Logging Connections to the ASA FirePOWER Module or External Server 36-4
Understanding How Access Control and SSL Rule Actions Affect Logging 36-4
License Requirements for Connection Logging 36-7
Logging Security Intelligence (Blacklisting) Decisions 36-8
Logging Connections Based on Access Control Handling 36-9
Logging Connections Matching an Access Control Rule 36-10
Logging Connections Handled by the Access Control Default Action 36-11
Logging URLs Detected in Connections 36-13
Logging Encrypted Connections 36-14
Logging Decryptable Connections with SSL Rules 36-14
Setting Default Logging for Encrypted and Undecryptable Connections 36-15

CHAPTER 37 **Viewing Events** 37-1
Accessing ASA FirePOWER Real-Time Events 37-1
Understanding ASA FirePOWER Event Types 37-2
Event Fields in ASA FirePOWER Events 37-3
Intrusion Rule Classifications 37-12

CHAPTER 38 **Configuring External Alerting** 38-1
Working with Alert Responses 38-2
Creating an SNMP Alert Response 38-2
Creating a Syslog Alert Response 38-3
Modifying an Alert Response 38-5
Deleting an Alert Response 38-6
Enabling and Disabling Alert Responses 38-6
38-6

CHAPTER 39 **Configuring External Alerting for Intrusion Rules** 39-1
Using SNMP Responses 39-1
Configuring SNMP Responses 39-3
Using Syslog Responses 39-4
Configuring Syslog Responses 39-6

CHAPTER 40 **Using the ASA FirePOWER Dashboard** 40-1
Understanding Dashboard Widgets 40-1
Understanding Widget Preferences 40-1
Understanding the Predefined Widgets 40-2
Understanding the Appliance Information Widget 40-2
Understanding the Current Interface Status Widget 40-3
Understanding the Disk Usage Widget 40-3
Understanding the Product Licensing Widget 40-4
Understanding the Product Updates Widget 40-4
Understanding the System Load Widget 40-5
Understanding the System Time Widget 40-5
Working with the Dashboard 40-5
Viewing the Dashboard 40-6
Modifying the Dashboard 40-6

CHAPTER 41 **Using ASA FirePOWER Reporting** 41-1
Understanding Available Reports 41-1
Report Basics 41-2
Understanding Report Data 41-3
Drilling into Reports 41-3
Changing the Report Time Range 41-3
Controlling the Data Displayed in Reports 41-4
Understanding Report Columns 41-5

CHAPTER 42 **Scheduling Tasks** 42-1
Configuring a Recurring Task 42-1
Automating Backup Jobs 42-2
Automating Applying an Intrusion Policy 42-3
Automating Geolocation Database Updates 42-4
Automating Software Updates 42-5
Automating Software Downloads 42-5
Automating Software Installs 42-6
Automating URL Filtering Updates 42-7
Viewing Tasks 42-8
Using the Calendar 42-8
Using the Task List 42-9
Editing Scheduled Tasks 42-9
Deleting Scheduled Tasks 42-10

Deleting a Recurring Task 42-10
Deleting a One-Time Task 42-11

CHAPTER 43 **Managing System Policies** 43-1
Creating a System Policy 43-1
Editing a System Policy 43-2
Applying a System Policy 43-2
Deleting System Policies 43-3
Configuring a System Policy 43-3
Configuring the Access List for Your Appliance 43-3
Configuring Audit Log Settings 43-5
Configuring a Mail Relay Host and Notification Address 43-6
Configuring SNMP Polling 43-8
Enabling STIG Compliance 43-9
43-10

CHAPTER 44 **Configuring ASA FirePOWER Module Settings** 44-1
Viewing and Modifying the Appliance Information 44-1
Enabling Cloud Communications 44-2
System Information 44-4
Time 44-4

CHAPTER 45 **Licensing the ASA FirePOWER Module** 45-1
Understanding Licensing 45-1
Viewing Your Licenses 45-4
Adding a License to the ASA FirePOWER module 45-4
Deleting a License 45-5
45-5

CHAPTER 46 **Updating ASA FirePOWER Module Software** 46-1
Understanding Update Types 46-1
Performing Software Updates 46-2
Planning for the Update 46-2
Understanding the Update Process 46-3
Updating the ASA FirePOWER Module Software 46-4
Monitoring the Status of Major Updates 46-6
Uninstalling Software Updates 46-7

Updating the Vulnerability Database 46-8
Importing Rule Updates and Local Rule Files 46-9
Using One-Time Rule Updates 46-10
Using Recurring Rule Updates 46-13
Importing Local Rule Files 46-14
Viewing the Rule Update Log 46-15
Updating the Geolocation Database 46-19

CHAPTER 47 **Monitoring the System** 47-1
Viewing Host Statistics 47-1
Monitoring System Status and Disk Space Usage 47-2
Viewing System Process Status 47-2
Understanding Running Processes 47-4
Understanding System Daemons 47-4
Understanding Executables and System Utilities 47-5

CHAPTER 48 **Using Backup and Restore** 48-1
Creating Backup Files 48-1
Creating Backup Profiles 48-3
Uploading Backups from a Local Host 48-4
Restoring the Appliance from a Backup File 48-4

APPENDIX A **Generating Troubleshooting Files** A-1

APPENDIX B **Importing and Exporting Configurations** B-1
Exporting Configurations B-1
Importing Configurations B-3

APPENDIX C **Viewing the Status of Long-Running Tasks** C-1
Viewing the Task Queue C-1
Managing the Task Queue C-2

APPENDIX D **Security, Internet Access, and Communication Ports** D-1
Internet Access Requirements D-1
Communication Ports Requirements D-2



Introduction to the Cisco ASA FirePOWER Module

The Cisco ASA FirePOWER module® is a module that can be deployed on Cisco ASA5506-X, ASA5506H-X, ASA5506W-X, ASA5508-X, ASA5512-X, ASA5515-X, ASA5516-X, ASA5525-X, ASA5545-X, ASA5555-X, ASA5585-X-SSP-10, ASA5585-X-SSP-20, ASA5585-X-SSP-40, ASA5585-X-SSP-60. The module is designed to help you handle network traffic in a way that complies with your organization's security policy—your guidelines for protecting your network. A security policy may also include an acceptable use policy (AUP), which provides employees with guidelines of how they may use your organization's systems.

This guide provides information about onbox configuration of the features and functionality of the ASA FirePOWER module, accessible via ASDM. The explanatory text, diagrams, and procedures in each chapter provide detailed information to help you navigate the user interface, maximize the performance of your system, and troubleshoot complications.

Note If you enable command authorization on the ASA that hosts the ASA FirePOWER 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.

The topics that follow introduce you to the ASA FirePOWER module, describe its key components, and help you understand how to use this guide:

- Introduction to the ASA FirePOWER Module, page 1-1
- ASA FirePOWER Module Components, page 1-2
- License Conventions, page 1-3
- IP Address Conventions, page 1-4

Introduction to the ASA FirePOWER Module

The ASA FirePOWER module runs on an ASA device installed on network segments monitor traffic for analysis.

Deployed inline, the system can affect the flow of traffic using *access control*, which allows you to specify, in a granular fashion, how to handle the traffic entering, exiting, and traversing your network. The data that you collect about your network traffic and all the information you glean from it can be used to filter and control that traffic based on:

- simple, easily-determined transport and network layer characteristics: source and destination, port, protocol, and so on
- the latest contextual information on the traffic, including characteristics such as reputation, risk, business relevance, application used, or URL visited
- Microsoft Active Directory LDAP users in your organization

Each type of traffic inspection and control occurs where it makes the most sense for maximum flexibility and performance. For example, reputation-based blacklisting, because it uses simple source and destination data, can block prohibited traffic early in the process, while detecting and blocking intrusions and exploits is a last-line defense.

ASA FirePOWER Module Components

The topics that follow describe some of the key capabilities of the ASA FirePOWER module that contribute to your organization's security, acceptable use policy, and traffic management strategy:

- Access Control, page 1-2
- Intrusion Detection and Prevention, page 1-2
- Advanced Malware Protection and File Control, page 1-3
- Application Programming Interfaces, page 1-3

Access Control

Access control is a policy-based feature that allows you to specify, inspect, and log the traffic that can traverse your network. An *access control policy* determines how the system handles traffic on your network.

The simplest access control policy handles all traffic using its *default action*. You can set this default action to block or trust all traffic without further inspection, or to inspect traffic for intrusions.

A more complex access control policy can blacklist traffic based on Security Intelligence data, as well as use *access control rules* to exert granular control over network traffic logging and handling. These rules can be simple or complex, matching and inspecting traffic using multiple criteria; you can control traffic by security zone, network or geographical location, port, application, requested URL, ISE attribute, and user. Advanced access control options include decryption, preprocessing, and performance.

Each access control rule also has an *action*, which determines whether you monitor, trust, block, or allow matching traffic. When you allow traffic, you can specify that the system first inspect it with intrusion or file policies to block any exploits, malware, or prohibited files before they reach your assets or exit your network.

Intrusion Detection and Prevention

Intrusion detection and prevention is the system's last line of defense before traffic is allowed to its destination. *Intrusion policies* are defined sets of intrusion detection and prevention configurations invoked by your access control policy. Using *intrusion rules* and other settings, these policies inspect traffic for security violations and, in inline deployments, can block or alter malicious traffic.

If the system-provided policies do not fully address the security needs of your organization, custom policies can improve the performance of the system in your environment and can provide a focused view of the malicious traffic and policy violations occurring on your network. By creating and tuning custom policies you can configure, at a very granular level, how the system processes and inspects the traffic on your network for intrusions.

Advanced Malware Protection and File Control

To help you identify and mitigate the effects of malware, the ASA FirePOWER module's file control and advanced malware protection components can detect, track, capture, analyze, and optionally block the transmission of files (including malware files and nested files inside archive files) in network traffic.

File Control

File control allows devices to detect and block your users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols. You configure file control as part of your overall access control configuration; file policies associated with access control rules inspect network traffic that meets rule conditions.

Network-Based Advanced Malware Protection (AMP)

Network-based *advanced malware protection* (AMP) allows the system to inspect network traffic for malware in several types of files.

Regardless of whether you store a detected file, you can submit it to the Collective Security Intelligence Cloud for a simple known-disposition lookup using the file's SHA-256 hash value. Using this contextual information, you can configure the system to block or allow specific files.

You configure malware protection as part of your overall access control configuration; file policies associated with access control rules inspect network traffic that meets rule conditions.

Application Programming Interfaces

There are several ways to interact with the system using application programming interfaces (APIs). For detailed information, you can download additional documentation from either of the following Support Sites:

- **Cisco:** (http://www.cisco.com/cisco/web/support/index.html)

License Conventions

The License statement at the beginning of a section indicates the license required to use the feature described in the section, as follows:

Protection

A Protection license allows devices to perform intrusion detection and prevention, file control, and Security Intelligence filtering.

Control

A Control license allows devices to perform user and application control. A Control license requires a Protection license.

URL Filtering

A URL Filtering license allows devices to use regularly updated cloud-based category and reputation data to determine which traffic can traverse your network, based on the URLs requested by monitored hosts. A URL Filtering license requires a Protection license.

Malware

A Malware license allows devices to perform network-based advanced malware protection (AMP), that is, to detect, capture, and block malware in files transmitted over your network. A Malware license requires a Protection license.

Because licensed capabilities are often additive, this documentation only provides the highest required license for each feature. For example, if a feature requires Protection and Control licenses, only Control is listed. However, if functionality requires licenses that are not additive, the documentation lists them with a plus (+) character.

An "or" statement in a License statement indicates that a particular license is required to use the feature described in the section, but an additional license can add functionality. For example, within a file policy, some file rule actions require a Protection license while others require a Malware license. So, the License statement for the documentation on file rules lists "Protection or Malware."

IP Address Conventions

You can use IPv4 Classless Inter-Domain Routing (CIDR) notation and the similar IPv6 prefix length notation to define address blocks in many places in the ASA FirePOWER module.

CIDR notation uses a network IP address combined with a bit mask to define the IP addresses in the specified block of addresses. For example, the following table lists the private IPv4 address spaces in CIDR notation.

Table 1-1 CIDR Notation Syntax Examples

CIDR Block	IP Addresses in CIDR Block	Subnet Mask	Number of IP Addresses
10.0.0.0/8	10.0.0.0 - 10.255.255.255	255.0.0.0	16,777,216
172.16.0.0/12	172.16.0.0 - 172.31.255.255	255.240.0.0	1,048,576
192.168.0.0/16	192.168.0.0 - 192.168.255.255	255.255.0.0	65,536

Similarly, IPv6 uses a network IP address combined with a prefix length to define the IP addresses in a specified block. For example, 2001:db8::/32 specifies the IPv6 addresses in the 2001:db8:: network with a prefix length of 32 bits, that is, 2001:db8:: through 2001:db8:ffff:ffff:ffff:ffff:ffff:ffff.

When you use CIDR or prefix length notation to specify a block of IP addresses, the ASA FirePOWER module uses **only** the portion of the network IP address specified by the mask or prefix length. For example, if you type 10.1.2.3/8, the ASA FirePOWER module uses 10.0.0.0/8.

In other words, although Cisco recommends the standard method of using a network IP address on the bit boundary when using CIDR or prefix length notation, the ASA FirePOWER module does not require it.



Managing Reusable Objects

For increased flexibility and ease-of-use, the ASA FirePOWER module allows you to create named *objects*, which are reusable configurations that associate a name with a value so that when you want to use that value, you can use the named object instead.

You can create the following types of objects:

- network-based objects that represent IP addresses and networks, port/protocol pairs, security zones, and origin/destination country (geolocation)
- objects that help you handle unencrypted and decrypted traffic, including Security Intelligence feeds and lists, application filters, URLs, file lists, and intrusion policy variable sets

You can use these objects in various places in the ASA FirePOWER module, including access control policies, network analysis policies, intrusion policies and rules, reports, dashboards, and so on.

Grouping objects allows you to reference multiple objects with a single configuration. You can group network, port, and URL, and public key infrastructure (PKI) objects.

Note In most cases, editing an object used in a policy requires redeploying your configuration for your changes to take effect.

For more information, see the following sections:

- Using the Object Manager, page 2-2
- Working with Network Objects, page 2-3
- Working with Security Intelligence Lists and Feeds, page 2-4
- Working with Port Objects, page 2-9
- Working with URL Objects, page 2-10
- Working with Application Filters, page 2-11
- Working with Variable Sets, page 2-13
- Working with Sinkhole Objects, page 2-27
- Working with File Lists, page 2-28
- Working with Security Zones, page 2-32
- Working with Cipher Suite Lists, page 2-32
- Working with Distinguished Name Objects, page 2-33
- Working with PKI Objects, page 2-35

- Working with Geolocation Objects, page 2-42
- Working with Security Group Tag Objects, page 2-43

Using the Object Manager

License: Any

Create and manage objects, including application filters, variable sets, and security zones, using the object manager (**Configuration > ASA FirePOWER Configuration > Object Management**). You can group network, port, and URL and PKI objects; you can also sort, filter, and browse the list of objects and object groups.

For more information, see:

- Grouping Objects, page 2-2
- Browsing, Sorting, and Filtering Objects, page 2-3

Grouping Objects

License: Any

You can group network, port, PKI, and URL objects. The system allows you to use objects and object groups interchangeably. For example, anywhere you would use a port object, you can also use a port object group. Objects and object groups of the same type cannot have the same name.

When you edit an object group used in a policy (for example, a network object group used in an access control policy), you must redeploy the configuration for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Deleting a group does not delete the objects in the group, just their association with each other. Additionally, you cannot delete a group that is in use. For example, you cannot delete a URL group that you are using in a URL condition in a saved access control policy.

To group reusable objects:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under the type of **Network**, **Port**, **URL**, **PKI,** or **Distinguished Name** object you want to group, choose **Object Groups**.

Step 3 Click the **Add** button that corresponds with the object you want to group.

Step 4 Enter a **Name** for the group. You can use any printable standard ASCII characters except curly braces ({}).

Step 5 Choose one or more objects and click **Add**.

- Use Shift and Ctrl to choose multiple objects, or right-click and **Select All**.
- Use the filter field ( ) to search for existing objects to include, which updates as you type to display matching items. Click the reload icon ( ) above the search field or click the clear icon ( ) in the search field to clear the search string.
- Click the add icon ( ) to create objects on the fly if no existing objects meet your needs.

Step 6 Click **Store ASA FirePOWER Changes**.

Browsing, Sorting, and Filtering Objects

License: Any

The object manager displays 20 objects or groups per page. If you have more than 20 of any type of object or group, use the navigation links at the bottom of the page to view additional pages. You can also go to a specific page or click the refresh icon ( ) to refresh your view.

By default, the page lists objects and groups alphabetically by name. However, you can sort each type of object or group by any column in the display. An up (▲) or down (▶) arrow next to a column heading indicates that the page is sorted by that column in that direction. You can also filter the objects on the page by name or value.

To sort objects or groups:

Step 1 Click a column heading. To sort in the opposite direction, click the heading again.

To filter objects or groups:

Step 1 Enter your filter criteria in the **Filter** field.

The page updates as you type to display matching items. The field accepts one or more asterisks (*) as wild cards.

Working with Network Objects

License: Any

A network object represents one or more IP addresses that you can specify either individually or as address blocks. You can use network objects and groups (see Grouping Objects, page 2-2) in various places in the ASA FirePOWER module, including access control policies, network variables, intrusion rules, reports, and so on.

You also cannot delete a network object that is in use. Additionally, after you edit a network object used in an access control or intrusion policy, you must redeploy policies for your changes to take effect.

To create a network object:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **Network**, choose **Individual Objects**.

Step 3 Click **Add Network**.

Step 4 Enter a **Name** for the network object. You can use any printable standard ASCII characters except curly braces ({}).

Step 5 For each IP address or address block you want to add to the network object, enter its value and click **Add**.

Step 6 Click **Store ASA FirePOWER Changes**.

Step 7 If an active policy references your object, deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Working with Security Intelligence Lists and Feeds

License: Protection

The Security Intelligence feature allows you to, per access control policy, specify the traffic that can traverse your network based on the source or destination IP address. This is especially useful if you want to blacklist — deny traffic to and from — specific IP addresses, before the traffic is subjected to analysis by access control rules. Similarly, you can add IP addresses to the whitelist to force the system to handle their connections using access control.

If you are not sure whether you want to blacklist a particular IP address, you can use a "monitor-only" setting, which allows the system to handle the connection using access control, but also logs the connection's match to the blacklist.

A *global whitelist* and *global blacklist* are included by default in every access control policy, and apply to any zone. Additionally, within each access control policy, you can build a separate whitelist and blacklist using a combination of network objects and groups as well as Security Intelligence lists and feeds, all of which you can constrain by security zone.

Comparing Feeds and Lists

A Security Intelligence *feed* is a dynamic collection of IP addresses that the system downloads from an HTTP or HTTPS server at the interval you configure. Because feeds are regularly updated, the system can use up-to-date information to filter your network traffic. To help you build blacklists, the ASA FirePOWER module provides the *Intelligence Feed*, which represents IP addresses determined by the VRT to have a poor reputation.

Although it may take a few minutes for a feed update to take effect, you do not have to deploy policies after you create or modify a feed, or after a scheduled feed update.

Note If you want strict control over when the system downloads a feed from the Internet, you can disable automatic updates for that feed. However, Cisco recommends that you allow automatic updates. Although you can manually perform on-demand updates, allowing the system to download feeds on a regular basis provides you with the most up-to-date, relevant data.

In contrast with a feed, a Security Intelligence *list* is a simple static list of IP addresses that you manually upload to the system. Use custom lists to augment and fine-tune feeds and the global whitelist and blacklist. Note that editing custom lists (as well as editing network objects and removing IP addresses from the global whitelist or blacklist) require you to redeploy the configuration for your changes to take effect.

Formatting and Corrupt Feed Data

Feed and list source must be a simple text file no larger than 500MB, with one IP address or address block per line. Comment lines must start with the `#` character. List source files must use the `.txt` extension.

If the system downloads a corrupt feed or a feed with no recognizable IP addresses, the system continues using the old feed data (unless it is the first download). However, if the system can recognize even one IP address in the feed, it updates the addresses it can recognize.

Internet Access and High Availability

The system uses port 443/HTTPS to download the Intelligence Feed, and either 443/HTTP or 80/HTTP to download custom or third-party feeds. To update feeds, you must open the appropriate port, both inbound and outbound, on the device. If your system does not have direct access to the feed site, it can use a proxy server.

Note The system does **not** perform peer SSL certificate verification when downloading custom feeds, nor does the system support the use of certificate bundles or self-signed certificates to verify the remote peer.

Managing Feeds and Lists

You create and manage Security Intelligence lists and feeds, collectively called Security Intelligence objects, using the object manager's Security Intelligence page.

Note that you cannot delete a custom list or feed that is currently being used in a saved or applied access control policy. You also cannot delete a global list, although you can remove individual IP addresses. Similarly, although you cannot delete the Intelligence Feed, editing it allows you to disable or change the frequency of its updates.

Security Intelligence Object Quick Reference

The following table provides a quick reference to the objects you can use to perform Security Intelligence filtering.

Table 2-1 Security Intelligence Object Capabilities

Capability	Global Whitelist or Blacklist	Intelligence Feed	Custom Feed	Custom List	Network Object
method of use	in access control policies by default	in any access control policy as either a whitelist or blacklist object			
can be constrained by security zone?	no	yes	yes	yes	yes
can be deleted?	no	no	yes, unless currently being used in a saved or applied access control policy		
object manager edit capabilities	delete IP addresses only	disable or change update frequency	fully modify	upload a modified list only	fully modify
requires configuration redeployment when modified?	yes when deleting (adding IP addresses does not require redeploy)	no	no	yes	yes

For more information on creating, managing, and using Security Intelligence lists and feeds, see:

- Working with the Global Whitelist and Blacklist, page 2-6
- Working with the Intelligence Feed, page 2-6
- Working with Custom Security Intelligence Feeds, page 2-7

- Manually Updating Security Intelligence Feeds, page 2-7
- Working with Custom Security Intelligence Lists, page 2-8
- Blacklisting Using Security Intelligence IP Address Reputation, page 5-1

Working with the Global Whitelist and Blacklist

License: Protection

The system's global whitelist and blacklist are included by default in every access control policy, and apply to any zone. You can opt not to use these global lists on a per-policy basis.

You do not have to redeploy your configuration after adding an IP address to a global list. Conversely, after you delete IP addresses from the global whitelist or blacklist, you must redeploy your configuration for your changes to take effect.

Note that although you can add network objects with a netmask of /0 to the whitelist or blacklist, address blocks using a `/0` netmask in those objects are ignored and whitelist and blacklist filtering does not occur based on those addresses. Address blocks with a `/0` netmask from security intelligence feeds is also ignored. If you want to monitor or block all traffic targeted by a policy, instead of security intelligence filtering, use an access control rule with the **Monitor** or **Block** rule action, respectively, and a default value of **any** for the **Source Networks** and **Destination Networks**.

To remove IP addresses from the global whitelist or blacklist:

Step 1 On the object manager's Security Intelligence page, next to the global whitelist or blacklist, click the edit icon ( ).

Step 2 Next to the IP addresses you want to remove from the list, click the delete icon ( ).

To delete multiple IP addresses at once, use the Shift and Ctrl keys to choose them, then right-click and choose **Delete**.

Step 3 Click **Store ASA FirePOWER Changes**.

Step 4 If an active policy references your object, deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Working with the Intelligence Feed

License: Protection

To help you build blacklists, the ASA FirePOWER module provides the Intelligence Feed, which is comprised of several regularly updated lists of IP addresses determined by the VRT to have a poor reputation. Each list in the feed represents a specific category: open relays, known attackers, bogus IP addresses (bogon), and so on. In an access control policy, you can blacklist any or all of the categories.

Because the intelligence feed is regularly updated, the system can use up-to-date information to filter your network traffic. Malicious IP addresses that represent security threats such as malware, spam, botnets, and phishing may appear and disappear faster than you can update and deploy new policies.

Although you cannot delete the Intelligence Feed, editing it allows you to change the frequency of its updates. By default, the feed updates every two hours.

To modify the intelligence feed's update frequency:

Step 1 On the object manager's Security Intelligence page, next to the Intelligence Feed, click the edit icon ( ).

Step 2 Edit the **Update Frequency**.

You can choose various intervals from two hours to one week. You can also disable feed updates.

Step 3 Click **Store ASA FirePOWER Changes**.

Working with Custom Security Intelligence Feeds

License: Protection

Custom or third-party Security Intelligence feeds allow you to augment the Intelligence Feed with other regularly-updated reputable whitelists and blacklists on the Internet. You can also set up an internal feed.

When you configure a feed, you specify its location using a URL; the URL cannot be Punycode-encoded. By default, the system downloads the entire feed source on the interval you configure.

Optionally, you can configure the system to use an md5 checksum to determine whether to download an updated feed. If the checksum has not changed since the last time the module downloaded the feed, the system does not need to re-download it. You may want to use md5 checksums for internal feeds, especially if they are large. The md5 checksum must be stored in a simple text file with only the checksum. Comments are not supported.

To configure a Security Intelligence feed:

Step 1 On the object manager's Security Intelligence page, click **Add Security Intelligence**.

Step 2 Enter a **Name** for the feed. You can use any printable standard ASCII characters except curly braces (`{}`).

Step 3 From the **Type** drop-down list, specify that you want to configure a **Feed**.

Step 4 Specify a **Feed URL** and optionally, an **MD5 URL**.

Step 5 Specify an **Update Frequency**.

You can choose various intervals from two hours to one week. You can also disable feed updates.

Step 6 Click **Store ASA FirePOWER Changes**.

The Security Intelligence feed object is created. Unless you disabled feed updates, the system attempts to download and verify the feed. You can now use the feed object in access control policies.

Manually Updating Security Intelligence Feeds

License: Protection

Manually updating Security Intelligence feeds updates all feeds, including the Intelligence Feed.

To update all Security Intelligence feeds:

Step 1 On the object manager's Security Intelligence page, click **Update Feeds**.

Step 2 Confirm that you want to update all feeds.

The system warns that it can take several minutes for the update to take effect.

Step 3 Click **OK**.

After the system downloads and verifies the feed updates, it begins filtering traffic using the updated feeds.

Working with Custom Security Intelligence Lists

License: Protection

A Security Intelligence list is a simple static list of IP addresses and address blocks that you manually upload. Custom lists are useful if you want to augment and fine-tune feeds or one of the global lists.

Note that netmasks for address blocks can be integers from `0` to `32` or `0` to `128`, for IPv4 and IPv6, respectively.

For example, if a reputable feed improperly blocks your access to vital resources but is overall useful to your organization, you can create a custom whitelist that contains only the improperly classified IP addresses, rather than removing the Security Intelligence feed object from the access control policy's blacklist.

Note that to modify a Security Intelligence list, you must make your changes to the source file and upload a new copy. For more information, see Updating a Security Intelligence List, page 2-8.

To upload a new Security Intelligence list:

Step 1 On the object manager's Security Intelligence page, click **Add Security Intelligence**.

Step 2 Enter a **Name** for the list. You can use any printable standard ASCII characters except curly braces (`{}`).

Step 3 From the **Type** drop-down list, specify that you want to upload a **List**.

Step 4 Click **Browse** to browse to the list`.txt` file, then click **Upload**.

The list is uploaded. The pop-up window displays the total number of IP addresses and address blocks that the system found in the list.

If the number is not what you expected, check the formatting of the file and try again.

Step 5 Click **Store ASA FirePOWER Changes**.

Updating a Security Intelligence List

License: Protection

To edit a Security Intelligence list, you must make your changes to the source file and upload a new copy. You cannot modify the file's contents using ASDM. If you do not have access to the source file, you can download a copy using the ASDM interface.

To modify a Security Intelligence list:

Step 1 On the object manager's Security Intelligence page, next to the list you want to update, click the edit icon ( ).

Step 2 If you need a copy of the list to edit, click **Download**, then follow the prompts to save the list as a text file.

Step 3 Make changes to the list as necessary.

Step 4 On the Security Intelligence pop-up window, click **Browse** to browse to the modified list, then click **Upload**.

Step 5 Click **Store ASA FirePOWER Changes**.

Step 6 If an active policy references your object, deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Working with Port Objects

License: Any

Port objects represent different protocols in slightly different ways:

- For TCP and UDP, a port object represents the transport layer protocol, with the protocol number in parentheses, plus an optional associated port or port range. For example: `TCP(6)/22`.
- For ICMP and ICMPv6 (IPv6-ICMP), the port object represents the internet layer protocol plus an optional type and code. For example: `ICMP(1):3:3`.
- A port object can also represent other protocols that do not use ports.

Note that the system provides default port objects for well-known ports. You can modify or delete these objects, but Cisco recommends that you create custom port objects instead.

You can use port objects and groups (see Grouping Objects, page 2-2) in various places in the ASA FirePOWER module, including access control policies and port variables.

You cannot delete a port object that is in use. Additionally, after you edit or delete a port object, if an active policy references the object, you must redeploy your configuration for the changes to take effect; see Deploying Configuration Changes, page 4-11.

Note that you cannot add any protocol other than TCP or UDP for source port conditions in access control rules. Also, you cannot mix transport protocols when setting both source and destination port conditions in a rule.

If you add an unsupported protocol to a port object group used in a source port condition, the rule where it is used does not apply on policy deploy. Additionally, if you create a port object containing both TCP and UDP ports, then add it as a source port condition in a rule, you cannot add a destination port, and vice versa.

To create a port object:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **Port**, choose **Individual Objects**.

Step 3 Click **Add Port**.

Step 4 Enter a **Name** for the port object. You can use any printable standard ASCII characters except curly braces (`{}`).

Step 5 Choose a **Protocol**.

You can quickly choose **TCP**, **UDP**, **IP**, **ICMP**, or **IPv6-ICMP**, or you can use the **Other** drop-down list to choose either a different protocol or **All** protocols.

Step 6 Optionally, restrict a TCP or UDP port object using a **Port** or port range.

You can specify any port from 1 to 65535 or `any` to match all ports. Use a hyphen to specify a range of ports.

Step 7 Optionally, restrict an ICMP or IPV6-ICMP port object using a **Type** and, if appropriate, a related **Code**.

When you create an ICMP or IPv6-ICMP object, you can specify the type and, if applicable, the code. For more information on ICMP types and codes, see http://www.iana.org/assignments/icmp-parameters/icmp-parameters.xml and http://www.iana.org/assignments/icmpv6-parameters/icmpv6-parameters.xml. You can set the type to any to match any type or set the code to any to match any code for the specified type.

Step 8 Optionally, choose **Other** and a protocol from the drop-down list. If you choose **All** protocols, enter a port number in the **Port** field.

Step 9 Click **Store ASA FirePOWER Changes**.

Working with URL Objects

License: Any

Each URL object you configure represents a single URL or IP address. You can use URL objects and groups (see Grouping Objects, page 2-2) in access control policies. For example, you could write an access control rule that blocks a specific URL.

Note that to block HTTPS traffic, you can enter the URL from the Secure Sockets Layer (SSL) certificate for the traffic. When entering a URL from a certificate, enter the domain name and omit subdomain information. (For example, type `example.com` rather than `www.example.com`.) If you block traffic based on the certificate URL, both HTTP and HTTPS traffic to that website are blocked.

You cannot delete a URL object that is in use. Additionally, after you edit or delete a URL object, if an active policy references the object, you must redeploy your configuration for the changes to take effect; see Deploying Configuration Changes, page 4-11.

To create a URL object:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **URL**, choose **Individual Objects**.

Step 3 Click **Add URL**.

Step 4 Enter a **Name** for the URL object. You can use any printable standard ASCII characters except curly braces (`{}`).

Step 5 Enter the **URL** or IP address for the URL object.

Step 6 Click **Store ASA FirePOWER Changes**.

Working with Application Filters

License: Any

When the ASA FirePOWER module analyzes IP traffic, it attempts to identify the commonly used applications on your network. Application awareness is crucial to performing application-based access control. The system is delivered with detectors for many applications, and Cisco frequently updates and adds additional detectors via system and vulnerability database (VDB) updates.

Application filters group applications according to criteria associated with the applications' risk, business relevance, type, categories, and tags. Using application filters allows you to quickly create application conditions for access control rules because you do not have to search for and add applications individually; for more information, see Matching Traffic with Application Filters, page 8-3.

Another advantage to using application filters is that you do not have to update access control rules that use filters when you modify or add new applications. For example, if you configure your access control policy to block all social networking applications, and a VDB update includes a new social networking application detector, the policy is updated when you update the VDB. Although you must redeploy the changed configuration before the system can block the new application, you do not have to update the access control rule that blocks the application.

If the system-provided application filters do not group applications according to your needs, you can create your own filters. User-defined filters can group and combine system-provided filters. For example, you could create a filter that would allow you to block all very high risk, low business relevance applications. You can also create a filter by manually specifying individual applications, although you should keep in mind those filters do **not** automatically update when you update the module software or the VDB.

As with system-provided application filters, you can use user-defined application filters in access control rules.

You use the object manager (**Configuration > ASA FirePOWER Configuration > Object Management**) to create and manage application filters. Note that you can also create an application filter on the fly while adding an application condition to an access control rule.

The Application Filters list contains the system-provided application filters that you can choose to build your own filter. You can constrain the filters that appear by using a search string; this is especially useful for categories and tags.

The Available Applications list contains the individual applications in the filters you select. You can also constrain the applications that appear by using a search string.

The system links multiple filters of the same filter type with an OR operation. Consider a scenario where the medium risk filter contains 100 applications and the high risk filter contains 50 applications. If you choose both filters, the system would display 150 available applications.

The system links different types of filters with an AND operation. For example, if you choose the medium and high risk filters and the medium and high business relevance filters, the system displays the applications that have medium or high risk, and also have medium or high business relevance.

Tip Click an information icon ( ) for more information about the associated application. To display additional information, click any of the Internet search links in the information pop-up.

After you determine the applications you want to add to the filter, you can add them either individually, or, if you chose an application filter, **All apps matching the filter**. You can add multiple filters and multiple applications, in any combination, as long as the total number of items in the Selected Applications and Filters list does not exceed 50.

After you create the application filter, it is listed on the Application Filters page of the object manager. The page displays the total number of conditions that comprise each filter.

For information on sorting and filtering the application filters that appear, see Using the Object Manager, page 2-2. Note that you cannot delete an application filter that is in use. Additionally, after you edit or delete an application filter object, if an active policy references the object, you must redeploy your configuration for the changes to take effect; see Deploying Configuration Changes, page 4-11.

To create an application filter:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Click **Application Filters**.

Step 3 Click **Add Application Filter**.

Step 4 Enter a **Name**. You can use any printable standard ASCII characters except curly braces ({}).

Step 5 Optionally, use system-provided filters in the **Application Filters** list to narrow the list of applications you want to add to the filter:

- Click the arrow next to each filter type to expand and collapse the list.
- Right-click a filter type and click **Check All** or **Uncheck All**. Note that the list indicates how many filters you have selected of each type.
- To narrow the filters that appear, enter a search string in the **Search by name** field; this is especially useful for categories and tags. To clear the search, click the clear icon (✕).
- To refresh the filters list and clear any selected filters, click the reload icon (↻).
- To clear all filters and search fields, click **Clear All Filters**.

The applications that match the filters you select appear in the Available Applications list. The list displays 100 applications at a time.

Step 6 Choose the applications that you want to add to the filter from the **Available Applications** list:

- Choose **All apps matching the filter** to add all the applications that meet the constraints you specified in the previous step.
- To narrow the individual applications that appear, enter a search string in the **Search by name** field. To clear the search, click the clear icon (✕).
- Use the paging icons at the bottom of the list to browse the list of individual available applications.
- Use Shift and Ctrl keys to choose multiple individual applications. Right-click to **Select All** currently displayed individual applications.
- To refresh the applications list and clear any selected applications, click the reload icon (↻).

You cannot choose individual applications and **All apps matching the filter** at the same time.

Step 7 Add the selected applications to the filter. You can click and drag, or you can click **Add to Rule**.

The result is the combination of:

- the selected Application Filters
- either the selected individual Available Applications, or **All apps matching the filter**

You can add up to 50 applications and filters to the filter. To delete an application or filter from the selected applications, click the appropriate delete icon (🗑). You can also select one or more applications and filters, or right click to **Select All**, then right-click to **Delete Selected**.

Step 8 Click **Store ASA FirePOWER Changes**.

Working with Variable Sets

License: Protection

Variables represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions, adaptive profiles, and dynamic rule states.

Tip Preprocessor rules can trigger events regardless of the hosts defined by network variables used in intrusion rules.

You use variable sets to manage, customize, and group your variables. You can use the default variable set provided by the ASA FirePOWER module or create your own custom sets. Within any set you can modify predefined default variables and add and modify user-defined variables.

Most of the shared object rules and standard text rules that the ASA FirePOWER module provides use predefined default variables to define networks and port numbers. For example, the majority of the rules use the variable `$HOME_NET` to specify the protected network and the variable `$EXTERNAL_NET` to specify the unprotected (or outside) network. In addition, specialized rules often use other predefined variables. For example, rules that detect exploits against web servers use the `$HTTP_SERVERS` and `$HTTP_PORTS` variables.

Rules are more effective when variables more accurately reflect your network environment. At a minimum, you should modify default variables in the default set as described in Optimizing Predefined Default Variables, page 2-13. By ensuring that a variable such as `$HOME_NET` correctly defines your network and `$HTTP_SERVERS` includes all web servers on your network, processing is optimized and all relevant systems are monitored for suspicious activity.

To use your variables, you link variable sets to intrusion policies associated with access control rules or with the default action of an access control policy. By default, the default variable set is linked to all intrusion policies used by access control policies.

See the following sections for more information:

- Optimizing Predefined Default Variables, page 2-13
- Understanding Variable Sets, page 2-16
- Managing Variable Sets, page 2-17
- Managing Variables, page 2-18
- Adding and Editing Variables, page 2-20
- Resetting Variables, page 2-25
- Linking Variable Sets to Intrusion Policies, page 2-26
- Understanding Advanced Variables, page 2-27

Optimizing Predefined Default Variables

License: Protection

By default, the ASA FirePOWER module provides a single default variable set, which is comprised of predefined default variables. The Vulnerability Research Team (VRT) uses rule updates to provide new and updated intrusion rules and other intrusion policy elements, including default variables. See Importing Rule Updates and Local Rule Files, page 46-9 for more information.

Because many intrusion rules provided by the ASA FirePOWER module use predefined default variables, you should set appropriate values for these variables. Depending on how you use variable sets to identify traffic on your network, you can modify the values for these default variables in any or all variable sets. See Adding and Editing Variables, page 2-20 for more information.

Caution Importing an access control or an intrusion policy overwrites existing default variables in the default variable set with the imported default variables. If your existing default variable set contains a custom variable not present in the imported default variable set, the unique variable is preserved. For more information, see Importing Configurations, page B-3.

The following table describes the variables provided by the ASA FirePOWER module and indicates which variables you typically would modify. For assistance determining how to tailor variables to your network, contact Professional Services or Support.

Table 2-2 Variables Provided by the ASA FirePower Module

Variable Name	Description	Modify?
`$AIM_SERVERS`	Defines known AOL Instant Messenger (AIM) servers, and is used in chat-based rules and rules that look for AIM exploits.	Not required.
`$DNS_SERVERS`	Defines Domain Name Service (DNS) servers. If you create a rule that affects DNS servers specifically, you can use the `$DNS_SERVERS` variable as a destination or source IP address.	Not required in current rule set.
`$EXTERNAL_NET`	Defines the network that the ASA FirePOWER module views as the unprotected network, and is used in many rules to define the external network.	Yes, you should adequately define `$HOME_NET` and then exclude `$HOME_NET` as the value for `$EXTERNAL_NET`.
`$FILE_DATA_PORTS`	Defines non-encrypted ports used in intrusion rules that detect files in a network stream.	Not required.
`$FTP_PORTS`	Defines the ports of FTP servers on your network, and is used for FTP server exploit rules.	Yes, if your FTP servers use ports other than the default ports (you can view the default ports in the module interface).
`$GTP_PORTS`	Defines the data channel ports where the packet decoder extracts the payload inside a GTP (General Packet Radio Service [GPRS] Tunneling Protocol) PDU.	Not required.
`$HOME_NET`	Defines the network that the associated intrusion policy monitors, and is used in many rules to define the internal network.	Yes, to include the IP addresses for your internal network.
`$HTTP_PORTS`	Defines the ports of web servers on your network, and is used for web server exploit rules.	Yes, if your web servers use ports other than the default ports (you can view the default ports in the module interface).
`$HTTP_SERVERS`	Defines the web servers on your network. Used in web server exploit rules.	Yes, if you run HTTP servers.

Table 2-2 Variables Provided by the ASA FirePower Module (continued)

Variable Name	Description	Modify?
$ORACLE_PORTS	Defines Oracle database server ports on your network, and is used in rules that scan for attacks on Oracle databases.	Yes, if you run Oracle servers.
$SHELLCODE_PORTS	Defines the ports you want the system to scan for shell code exploits, and is used in rules that detect exploits that use shell code.	Not required.
$SIP_PORTS	Defines the ports of SIP servers on your network, and is used for SIP exploit rules.	Not required.
$SIP_SERVERS	Defines SIP servers on your network, and is used in rules that address SIP-targeted exploits.	Yes, if you run SIP servers, you should adequately define $HOME_NET and then include $HOME_NET as the value for $SIP_SERVERS.
$SMTP_SERVERS	Defines SMTP servers on your network, and is used in rules that address exploits that target mail servers.	Yes, if you run SMTP servers.
$SNMP_SERVERS	Defines SNMP servers on your network, and is used in rules that scan for attacks on SNMP servers.	Yes, if you run SNMP servers.
$SNORT_BPF	Identifies a legacy advanced variable that appears only when it existed on your system in a ASA FirePOWER module software release before Version 5.3.0 that you subsequently upgraded to Version 5.3.0 or greater. See Understanding Advanced Variables, page 2-27.	No, you can only view or delete this variable. You cannot edit it or recover it after deleting it.
$SQL_SERVERS	Defines database servers on your network, and is used in rules that address database-targeted exploits.	Yes, if you run SQL servers.
$SSH_PORTS	Defines the ports of SSH servers on your network, and is used for SSH server exploit rules.	Yes, if your SSH servers use ports other than the default port (you can view the default ports in the module interface).
$SSH_SERVERS	Defines SSH servers on your network, and is used in rules that address SSH-targeted exploits.	Yes, if you run SSH servers, you should adequately define $HOME_NET and then include $HOME_NET as the value for $SSH_SERVERS.
$TELNET_SERVERS	Defines known Telnet servers on your network, and is used in rules that address Telnet server-targeted exploits.	Yes, if you run Telnet servers.
$USER_CONF	Provides a general tool that allows you to configure one or more features not otherwise available via the module interface. See Understanding Advanced Variables, page 2-27. **Caution** Conflicting or duplicate $USER_CONF configurations will halt the system. See Understanding Advanced Variables, page 2-27.	No, only as instructed in a feature description or with the guidance of Support.

Understanding Variable Sets

License: Protection

Adding a variable to any set adds it to all sets; that is, each variable set is a collection of all variables currently configured on your system. Within any variable set, you can add user-defined variables and customize the value of any variable.

Initially, the ASA FirePOWER module provides a single, default variable set comprised of predefined default values. Each variable in the default set is initially set to its default value, which for a predefined variable is the value set by the VRT and provided in rule updates.

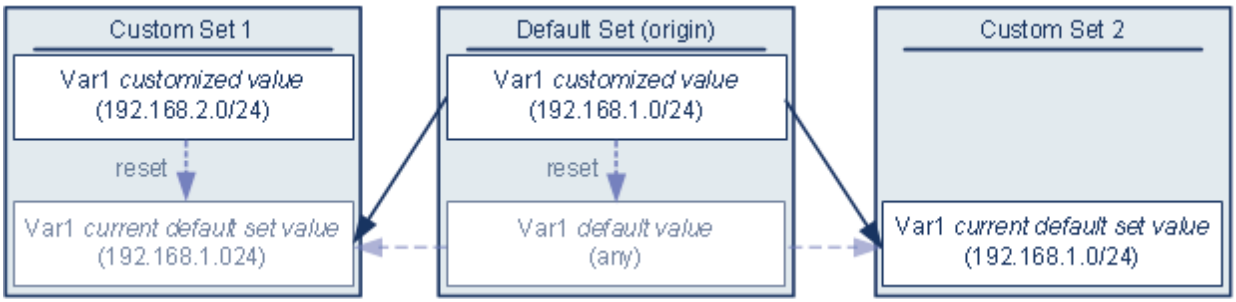
Although you can leave predefined default variables configured to their default values, Cisco recommends that you modify a subset of predefined variables as described in Optimizing Predefined Default Variables, page 2-13.

You could work with variables only in the default set, but in many cases you can benefit most by adding one or more custom sets, configuring different variable values in different sets, and perhaps even adding new variables.

When using multiple sets, it is important to remember that the *current value* of any variable in the default set determines the *default value* of the variable in all other sets.

Example: Adding a User-Defined Variable to the Default Set

The following diagram illustrates set interactions when you add the user-defined variable `Var1` to the default set with the value `192.168.1.0/24`.



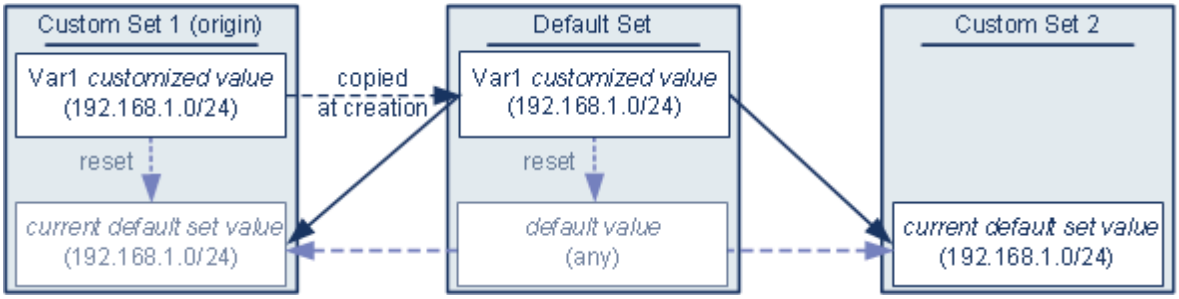
Optionally, you can customize the value of `Var1` in any set. In Custom Set 2 where `Var1` has not been customized, its value is `192.168.1.0/24`. In Custom Set 1 the customized value `192.168.2.0/24` of `Var1` overrides the default value. Resetting a user-defined variable in the default set resets its default value to `any` in all sets.

It is important to note in this example that, if you do not update `Var1` in Custom Set 2, further customizing or resetting `Var1` in the default set consequently updates the current, default value of `Var1` in Custom Set 2, thereby affecting any intrusion policy linked to the variable set.

Although not shown in the example, note that interactions between sets are the same for user-defined variables and default variables except that resetting a default variable in the default set resets it to the value configured by the system in the current rule update.

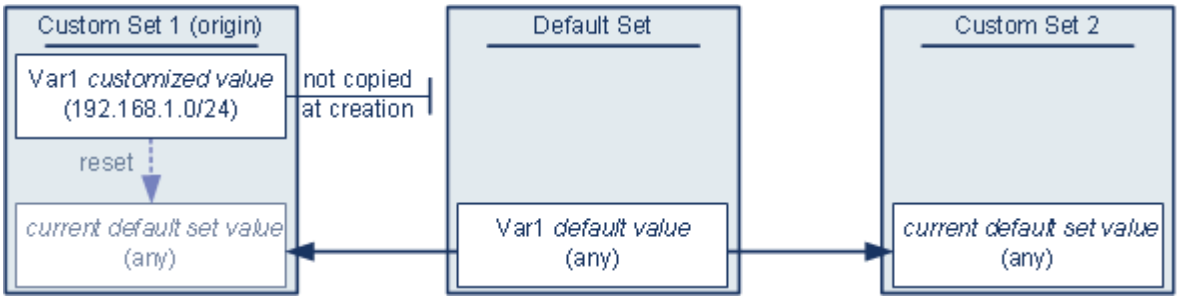
Examples: Adding a User-Defined Variable to a Custom Set

The next two examples illustrate variable set interactions when you add a user-defined variable to a custom set. When you save the new variable, you are prompted whether to use the configured value as the default value for other sets. In the following example, you elect **to use** the configured value.



Note that, except for the origin of `Var1` from Custom Set 1, this example is identical to the example above where you added `Var1` to the default set. Adding the customized value `192.168.1.0/24` for `Var1` to Custom Set 1 copies the value to the default set as a customized value with a default value of `any`. Thereafter, `Var1` values and interactions are the same as if you had added `Var1` to the default set. As with the previous example, keep in mind that further customizing or resetting `Var1` in the default set consequently updates the current, default value of `Var1` in Custom Set 2, thereby affecting any intrusion policy linked to the variable set.

In the next example, you add `Var1` with the value 192.168.1.0/24 to Custom Set 1 as in the previous example, but you elect **not to use** the configured value of `Var1` as the default value in other sets.



This approach adds `Var1` to all sets with a default value of `any`. After adding `Var1`, you can customize its value in any set. An advantage of this approach is that, by not initially customizing `Var1` in the default set, you decrease your risk of customizing the value in the default set and thus inadvertently changing the current value in a set such as Custom Set 2 where you have not customized `Var1`.

Managing Variable Sets

License: Protection

When you choose **Variable Sets** on the Object Manager page (**Configuration > ASA FirePOWER Configuration > Object Management**), the object manager lists the default variable set and any custom sets you created.

On a freshly installed system, the default variable set is comprised only of the default system-provided variables.

Each variable set includes the system-provided default variables and all custom variables you have added from any variable set. Note that you can edit the default set, but you cannot rename or delete the default set.

The following table summarizes the actions you can take to manage your variable sets.

Table 2-3 Variable Set Management Actions

To...	You can...
display your variable sets	choose **Configuration > ASA FirePOWER Configuration > Object Management**, then choose **Variable Set**.
filter variable sets by name	begin entering a name; as you type, the page refreshes to display matching names.
clear name filtering	click the clear icon (✖) in the filter field.
add a custom variable set	click **Add Variable Set**. For your convenience, new variable sets contain all currently defined default and customized variables.
edit a variable set	click the edit icon ( ) next to the variable set you want to edit. **Tip** You can also right-click within the row for a variable set, then choose **Edit**.
delete a custom variable set	click the delete icon ( ) next to the variable set, then click **Yes**. You cannot delete the default variable set. Note that variables created in a variable set you delete are not deleted or otherwise affected in other sets. **Tip** You can also right-click within the row for a variable set, choose **Delete**, then click **Yes**. Use the Ctrl and Shift keys to choose multiple sets.

After you configure variable sets, you can link them to intrusion policies.

To create or edit a variable set:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Choose **Variable Set**.

Step 3 Create a variable set or edit an existing set:

- To create a variable set, click **Add Variable Set**.
- To create a variable set, click the edit icon ( ) next to the variable set.

See Adding and Editing Variables, page 2-20 for information on adding and editing variables within a variable set.

Step 4 If an active policy references your object, deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Managing Variables

License: Protection

You manage variables on the new or edit variables page within a variable set. The variables page for all variable sets separates variables into Customized Variables and Default Variables page areas.

A *default variable* is a variable provided by the ASA FirePOWER module. You can customize the value of a default variable. You cannot rename or delete a default variable, and you cannot change its default value.

A *customized variable* is one of the following:

- customized default variables

  When you edit the value for a default variable, the system moves the variable from the Default Variables area to the Customized Variables area. Because variable values in the default set determine the default values of variables in custom sets, customizing a default variable in the default set modifies the default value of the variable in all other sets.

- user-defined variables

  You can add and delete your own variables, customize their values within different variable sets, and reset customized variables to their default values. When you reset a user-defined variable, it remains in the Customized Variables area.

The following table summarizes the actions you can take to create or edit variables.

Table 2-4 Variable Management Actions

To...	You can...
display the variables page	on the variable sets page, click **Add Variable Set** to create a new variable set, or click the edit icon ( ) next to the variable set you want to edit.
name and, optionally, describe your variable set	enter an alphanumeric string including spaces and special characters in the **Name** and **Description** fields.
add a variable	click **Add**. See Adding and Editing Variables, page 2-20 for more information.
edit a variable	click the edit icon ( ) next to the variable you want to edit. See Adding and Editing Variables, page 2-20 for more information.
reset a modified variable to its default value	click the reset icon ( ) next to a modified variable. A shaded reset icon indicates that the current value is already the default value.
delete a user-defined customized variable	click the delete icon ( ) next to the variable set; if you have saved the variable set since adding the variable, then click **Yes** to confirm that you want to delete the variable. You cannot delete default variables, and you cannot delete user-defined variables that are used by intrusion rules or other variables.
save changes to a variable set	click **Store ASA FirePOWER Changes**, then click **Yes** if the variable set is in use by an access control policy to confirm that you want to save your changes. Because the current value in the default set determines the default value in all other sets, modifying or resetting a variable in the default set changes the current value in other sets where you have not customized the default value.

To view the variables in a variable set:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Choose **Variable Set**.

Step 3 Create a variable set or edit an existing set:

- To create a variable set, click **Add Variable Set**.
- To create a variable set, click the edit icon ( ) next to the variable set.

Step 4 Create a variable or edit an existing variable:

- To create a variable, click **Add**.

- To edit a variable, click the edit icon ( ) next to the variable.

See Adding and Editing Variables, page 2-20 for information on adding and editing variables within a variable set.

Adding and Editing Variables

License: Protection

You can modify variables in any custom set.

If you create custom standard text rules, you might also want to create your own user-defined variables to more accurately reflect your traffic or as shortcuts to simplify the rule creation process. For example, if you create a rule that you want to inspect traffic in the "demilitarized zone" (or DMZ) only, you can create a variable named `$DMZ` whose value lists the server IP addresses that are exposed. You can then use the `$DMZ` variable in any rule written for this zone.

Adding a variable to a variable set adds it to all other sets. With one exception as explained below, the variable is added to other sets as the default value, which you can then customize.

When you add a variable from a custom set, you must choose whether to use the configured value as the customized value in the default set.

- If you **do use** the configured value (for example, 192.168.0.0/16), the variable is added to the default set using the configured value as a customized value with a default value of `any`. Because the current value in the default set determines the default value in other sets, the initial, default value in other custom sets is the configured value (which in the example is 192.168.0.0/16).
- If you **do not use** the configured value, the variable is added to the default set using only the default value `any` and, consequently, the initial, default value in other custom sets is `any`.

See Understanding Variable Sets, page 2-16 for more information.

You add variables within a variable set on the New Variable page and edit existing variables on the Edit Variable page. You use the two pages identically except that when you edit an existing variable you cannot change the variable name or variable type.

Each page consists mainly of three windows:

- available items, including existing network or port variables, objects, and network object groups
- networks or ports to include in the variable definition
- networks or ports to exclude from the variable definition

You can create or edit two types of variables:

- *network* variables specify the IP addresses of hosts in your network traffic. See Working with Network Variables, page 2-23.
- *port* variables specify TCP or UDP ports in network traffic, including the value `any` for either type. See Working with Port Variables, page 2-24.

When you specify whether you want to add a network or port variable type, the page refreshes to list available items. A search field above the list allows you to constrain the list, which updates as you type.

You can select and drag available items the list of items to include or exclude. You can also select items and click the **Include** or **Exclude** button. Use the Ctrl and Shift keys to choose multiple items. You can use the configuration field below the list of included or excluded items to specify literal IP addresses and address blocks for network variables, and ports and port ranges for port variables.

A list of items to include or exclude can be comprised of any combination of literal strings and existing variables, objects, and network object groups in the case of network variables.

The following table summarizes the actions you can take to create or edit your variables.

Table 2-5 Variable Edit Actions

To...	You can...
display the variables page	on the variable sets page, click **Add** to add a new variable, or click the edit icon ( ) next to an existing variable.
name your variable	in the **Name** field, enter a unique, case-sensitive alphanumeric string that includes no special characters other than the underscore character (_). Note that variable names are case-sensitive; for example, `var` and `Var` are each unique.
specify a network or port variable	choose **Network** or **Port** from the **Type** drop-down list. See Working with Network Variables, page 2-23 and Working with Port Variables, page 2-24 for detailed information on how you can use and configure network and port variables.
add an individual network object so you can then choose it from the list of available networks	choose **Network** from the **Type** drop-down list, then click the add icon ( ). See Working with Network Objects, page 2-3 for information on adding network objects using the object manager.
add an individual port object so you can then choose it from the list of available ports	choose **Port** from the **Type** drop-down list, then click the add icon ( ). Although you can add any port type, only TCP and UDP ports, including the value `any` for either type, are valid variable values, and the list of available ports only displays variables that use these value types. See Working with Port Objects, page 2-9 for information on adding port objects using the object manager.
search for available port or network items by name	begin entering a name in the search field above the list of available items; as you type, the page refreshes to display matching names.
clear name searching	click the reload icon ( ) above the search field or the clear icon ( ) in the search field.
differentiate between available items	look for items next to the variables icon ( ), network object icon ( ), port icon ( ), and object group icon ( ). Note that only network groups, not port groups, are available.
choose objects to include or exclude in the variable definition	click the object in the list of available networks or ports; use the Ctrl and Shift keys to choose multiple objects.
add selected items to the list of included or excluded networks or ports	drag and drop selected items. Alternately, click **Include** or **Exclude**. You can add network and port variables and objects from the list of available items. You can also add network object groups.
add a literal network or port to the list of networks or ports to include or exclude	click to remove the prompt from the literal **Network** or **Port** field, enter the literal IP address or address block for network variables, or the literal port or port range for port variables, then click **Add**. Note that you cannot enter domain names or lists; to add multiple items, add each individually.
add a variable with the value `any`	name the variable and specify the variable type, then click **Store ASA FirePOWER Changes** without configuring a value.

Table 2-5 Variable Edit Actions (continued)

To...	You can...
delete a variable or object from the included or excluded list	click the delete icon ( ) next to the variable.
save a new or modified variable	click **Store ASA FirePOWER Changes**; if you are adding a variable from custom set, then click **Yes** to use the configured value as the default value in other sets, or **No** to use a default value of `any`.

After you edit a variable, if an active policy references the object, you must redeploy your configuration for the changes to take effect; see Deploying Configuration Changes, page 4-11.

See the following sections for more information:

- Working with Network Variables, page 2-23
- Working with Port Variables, page 2-24

To create or edit a variable:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Choose **Variable Set**.

Step 3 Create a variable set or edit an existing set:

- To create a variable set, click **Add Variable Set**.
- To edit an existing variable set, click the edit icon ( ) next to the variable set.

Step 4 Create a new variable or edit an existing variable:

- To create a new variable, click **Add**.
- To edit an existing variable, click the edit icon ( ) next to the variable.

Step 5 If you are creating a new variable:

- Enter a unique variable **Name**.

  You can use alphanumeric characters and the underscore (_) character.
- Choose the **Network** or **Port** variable **Type** from the drop-down list.

Step 6 Optionally, move items from the list of available networks or ports to the list of included or excluded items.

You can choose one or more items and then drag and drop, or click **Include** or **Exclude**. Use the Ctrl and Shift keys to choose multiple items.

Tip If addresses or ports in the included and excluded lists for a network or port variable overlap, excluded addresses or ports take precedence.

Step 7 Optionally, enter a single literal value, then click **Add**.

For network variables, you can enter a single IP address or address block. For port variables you can add a single port or port range, separating the upper and lower values with a hyphen (-).

Repeat this step as needed to enter multiple literal values.

Step 8 Click **Store ASA FirePOWER Changes** to save the variable. If you are adding a new variable from a custom set, you have the following options:

- Click **Yes** to add the variable using the configured value as the customized value in the default set and, consequently, the default value in other custom sets.
- Click **No** to add the variable as the default value of `any` in the default set and, consequently, in other custom sets.

Step 9 When you have finished making changes, click **Store ASA FirePOWER Changes** to save the variable set, then click **Yes**.

Step 10 If an active policy references your object, deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Working with Network Variables

License: Protection

Network variables represent IP addresses you can use in intrusion rules that you enable in an intrusion policy and in intrusion policy rule suppressions, dynamic rule states, and adaptive profiles. Network variables differ from network objects and network object groups in that network variables are specific to intrusion policies and intrusion rules, whereas you can use network objects and groups to represent IP addresses in various places in the ASA FirePOWER module, including access control policies, network variables, intrusion rules, reports, and so on. See Working with Network Objects, page 2-3 for more information.

You can use network variables in the following configurations to specify the IP addresses of hosts on your network:

- intrusion rules

  Intrusion rule **Source IPs** and **Destination IPs** header fields allow you to restrict packet inspection to the packets originating from or destined to specific IP addresses. See Specifying IP Addresses In Intrusion Rules, page 30-5.

- suppressions

  The **Network** field in source or destination intrusion rule suppressions allows you to suppress intrusion event notifications when a specific IP address or range of IP addresses triggers an intrusion rule or preprocessor. See Configuring Suppression Per Intrusion Policy, page 27-25.

- dynamic rule states

  The **Network** field in source or destination dynamic rule states allows you to detect when too many matches for an intrusion rule or preprocessor rule occur in a given time period. See Adding Dynamic Rule States, page 27-28.

- adaptive profiles

  The adaptive profiles **Networks** field identifies hosts in the network where you want to improve reassembly of packet fragments and TCP streams in passive deployments. See Tuning Preprocessing in Passive Deployments, page 25-1.

When you use variables in the fields identified in this section, the variable set you link to an intrusion policy determines the variable values in the network traffic handled by an access control policy that uses the intrusion policy.

You can add any combination of the following network configurations to a variable:

- any combination of network variables, network objects, and network object groups that you select from the list of available networks

  See Working with Network Objects, page 2-3 for information on creating individual and group network objects using the object manager.

- individual network objects that you add from the New Variable or Edit Variable page, and can then add to your variable and to other existing and future variables

- literal, single IP addresses or address blocks

  You can list multiple literal IP addresses and address blocks by adding each individually. You can list IPv4 and IPv6 addresses and address blocks alone or in any combination. When specifying IPv6 addresses, you can use any addressing convention defined in RFC 4291.

The default value for included networks in any variable you add is the word `any`, which indicates any IPv4 or IPv6 address. The default value for excluded networks is none, which indicates no network. You can also specify the address `::` in a literal value to indicate any IPv6 address in the list of included networks, or no IPv6 addresses in the list of exclusions.

Adding networks to the excluded list negates the specified addresses and address blocks. That is, you can match any IP address with the exception of the excluded IP address or address blocks.

For example, excluding the literal address `192.168.1.1` specifies any IP address other than 192.168.1.1, and excluding `2001:db8:ca2e::fa4c` specifies any IP address other than 2001:db8:ca2e::fa4c.

You can exclude any combination of networks using literal or available networks. For example, excluding the literal values `192.168.1.1` and `192.168.1.5` *includes* any IP address other than 192.168.1.1 or 192.168.1.5. That is, the system interprets this as "**not** 192.168.1.1 **and not** 192.168.1.5," which matches any IP address other than those listed between brackets.

Note the following points when adding or editing network variables:

- You cannot logically exclude the value `any` which, if excluded, would indicate no address. For example, you cannot add a variable with the value `any` to the list of excluded networks.
- Network variables identify traffic for the specified intrusion rule and intrusion policy features. Note that preprocessor rules can trigger events regardless of the hosts defined by network variables used in intrusion rules.
- Excluded values must resolve to a subset of included values. For example, you cannot include the address block 192.168.5.0/24 and exclude 192.168.6.0/24. An error message warns you and identifies the offending variable, and you cannot save your variable set when you exclude a value outside the range of included values.

For information on adding and editing network variables, see Adding and Editing Variables, page 2-20.

Working with Port Variables

License: Protection

Port variables represent TCP and UDP ports you can use in the **Source Port** and **Destination Port** header fields in intrusion rules that you enable in an intrusion policy. Port variables differ from port objects and port object groups in that port variables are specific to intrusion rules. You can create port objects for protocols other than TCP and UDP, and you can use port objects in port variables and access control policies. See Working with Port Objects, page 2-9 for more information.

You can use port variables in the intrusion rule **Source Port** and **Destination Port** header fields to restrict packet inspection to packets originating from or destined to specific TCP or UDP ports.

When you use variables in these fields, the variable set you link to the intrusion policy associated with an access control rule or policy determines the values for these variables in the network traffic where the system applies the access control policy.

You can add any combination of the following port configurations to a variable:

- any combination of port variables and port objects that you cboose from the list of available ports

  Note that the list of available ports does not display port object groups, and you cannot add these to variables. See Working with Port Objects, page 2-9 for information on creating port objects using the object manager.

- individual port objects that you add from the New Variable or Edit Variable page, and can then add to your variable and to other existing and future variables

  Only TCP and UDP ports, including the value `any` for either type, are valid variable values. If you use the new or edit variables page to add a valid port object that is not a valid variable value, the object is added to the system but is not displayed in the list of available objects. When you use the object manager to edit a port object that is used in a variable, you can only change its value to a valid variable value.

- single, literal port values and port ranges

  You must separate port ranges with a dash (-). Port ranges indicated with a colon (:) are supported for backward compatibility, but you cannot use a colon in port variables that you create.

  You can list multiple literal port values and ranges by adding each individually in any combination.

Note the following points when adding or editing port variables:

- The default value for included ports in any variable you add is the word `any`, which indicates any port or port range. The default value for excluded ports is `none`, which indicates no ports.

Tip To create a variable with the value `any`, name and save the variable without adding a specific value.

- You cannot logically exclude the value `any` which, if excluded, would indicate no ports. For example, you cannot save a variable set when you add a variable with the value `any` to the list of excluded ports.
- Adding ports to the excluded list negates the specified ports and port ranges. That is, you can match any port with the exception of the excluded ports or port ranges.
- Excluded values must resolve to a subset of included values. For example, you cannot include the port range 10-50 and exclude port 60. An error message warns you and identifies the offending variable, and you cannot save your variable set when you exclude a value outside the range of included values.

For information on adding and editing port variables, see Adding and Editing Variables, page 2-20.

Resetting Variables

License: Protection

You can reset a variable to its default value on the variable set new or edit variables page. The following table summarizes the basic principles of resetting variables.

Table 2-6 Variable Reset Values

Resetting this variable type...	In this set type...	Resets it to...
default	default	the rule update value
user-defined	default	`any`
default or user-defined	custom	the current default set value (modified or unmodified)

Resetting a variable in a custom set simply resets it to the current value for that variable in the default set.

Conversely, resetting or modifying the value of a variable in the default set always updates the default value of that variable in all custom sets. When the reset icon is grayed out, indicating that you cannot reset the variable, this means that the variable has no customized value in that set. Unless you have customized the value for a variable in a custom set, a change to the variable in the default set updates the value used in any intrusion policy where you have linked the variable set.

Note It is good practice when you modify a variable in the default set to assess how the change affects any intrusion policy that uses the variable in a linked custom set, especially when you have not customized the variable value in the custom set.

When the customized value and the reset value are the same, this indicates one of the following:

- you are in the custom or default set where you added the variable with the value `any`
- you are in the custom set where you added the variable with an explicit value and elected to use the configured value as the default value

Linking Variable Sets to Intrusion Policies

License: Control

By default, the ASA FirePOWER module links the default variable set to all intrusion policies used in an access control policy. When you deploy an access control policy that uses an intrusion policy, intrusion rules that you have enabled in the intrusion policy use the variable values in the linked variable set.

When you modify a custom variable set used by an intrusion policy in an access control policy, the system reflects the status for that policy as out-of-date on the Access Control page. You must deploy the configuration to implement changes in your variable set. When you modify the default set, the system reflects the status of all access control policies that use intrusion policies as out-of-date, and you must redeploy the configuration to implement your changes.

See the following sections for information:

- To link a variable set other than the default set to an access control rule, see the procedure in Configuring an Access Control Rule to Perform Intrusion Prevention, page 11-4.
- To link a variable set other than the default set to the default action of an access control policy, see Setting Default Handling and Inspection for Network Traffic, page 4-4.
- To deploy access control policies, including policies that link variable sets to intrusion policies, see Deploying Configuration Changes, page 4-11.

Understanding Advanced Variables

License: Protection

Advanced variables allow you to configure features that you cannot otherwise configure via the module interface. The ASA FirePOWER module currently provides only two advanced variables, and you can only edit the USER_CONF advanced variable.

USER_CONF

USER_CONF provides a general tool that allows you to configure one or more features not otherwise available via the module interface.

Caution Do **not** use the advanced variable USER_CONF to configure an intrusion policy feature unless you are instructed to do so in the feature description or by Support. Conflicting or duplicate configurations will halt the system.

When editing USER_CONF, you can type up to 4096 total characters on a single line; the line wraps automatically. You can include any number of valid instructions or lines until you reach the 8192 maximum character length for a variable or a physical limit such as disk space. Use the backslash (\) line continuation character after any complete argument in a command directive.

Resetting USER_CONF empties it.

Working with Sinkhole Objects

License: Protection

A sinkhole object represents either a DNS server that gives non-routeable addresses for all domain names within the sinkhole, or an IP address that does not resolve to a server. You can reference the sinkhole object within a DNS policy rule to redirect matching traffic to the sinkhole. You must assign the object both an IPv4 address and an IPv6 address.

You cannot delete a sinkhole object that is in use. Additionally, after you edit a sinkhole object used in a DNS policy, you must redeploy the configuration for your changes to take effect; see Deploying Configuration Changes, page 4-11.

To create a sinkhole object:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Choose **Sinkhole** from the list of object types.

Step 3 Click **Add Sinkhole**.

Step 4 Enter a **Name**.

Step 5 Enter the **IPv4 Address** and **IPv6 Address** of your sinkhole.

Step 6 You have the following options:

- If you want to redirect traffic to a sinkhole server, choose **Log Connections to Sinkhole**.
- If you want to redirect traffic to a non-resolving IP address, choose **Block and Log Connections to Sinkhole**.

Step 7 If you want to assign an Indication of Compromise (IoC) type to your sinkhole, choose one from the **Type** drop-down.

Step 8 Click **Store ASA FirePOWER Changes**.

Working with File Lists

License: Malware

If you use network-based advanced malware protection (AMP), and the Collective Security Intelligence Cloud incorrectly identifies a file's disposition, you can add the file to a *file list* using a SHA-256 hash value to better detect the file in the future. Depending on the type of file list, you can do the following:

- To treat a file as if the cloud assigned a clean disposition, add the file to the *clean list*.
- To treat a file as if the cloud assigned a malware disposition, add the file to the *custom detection list.*

Because you manually specify the blocking behavior for these files, the system does not perform malware cloud lookups, even if the files are otherwise identified as malware by the cloud. Note that you must configure a rule in the file policy with either a **Malware Cloud Lookup** or **Block Malware** action and a matching file type to calculate a file's SHA value. For more information, see Working with File Rules, page 35-10.

The system's clean list and custom detection list are included by default in every file policy. You can opt not to use either or both lists on a per-policy basis.

Caution Do **not** include files on this list that are actually malware. The system does not block them, even if the cloud assigned the file's a Malware disposition, or if you added the file to the custom detection list.

Each file list can contain up to 10000 unique SHA-256 values. To add files to the file list, you can:

- upload a file so the system calculates and adds the file's SHA-256 value.
- enter a file's SHA-256 value directly.
- create and upload a comma-separated value (CSV) source file containing multiple SHA-256 values. All non-duplicate SHA-256 values are added to the file list.

When you add a file to a file list, edit a SHA-256 value in the file list, or delete SHA-256 values from the file list, you must redeploy the configuration for your changes to take effect; see Deploying Configuration Changes, page 4-11.

For more information on using file lists, see the following topics:

- Uploading Multiple SHA-256 Values to a File List, page 2-28
- Uploading an Individual File to a File List, page 2-30
- Adding a SHA-256 Value to the File List, page 2-30
- Modifying Files on a File List, page 2-31
- Downloading a Source File from a File List, page 2-31

Uploading Multiple SHA-256 Values to a File List

License: Malware

You can add multiple SHA-256 values to a file list by uploading a comma-separated value (CSV) source file containing a list of SHA-256 values and descriptions. The system validates the contents and populates the file list with valid SHA-256 values.

The source file must be a simple text file with a .csv file name extension. Any header must start with a pound sign (`#`); it is treated as a comment and not uploaded. Each entry should contain a single SHA-256 value followed by a description of up to 256 alphanumeric or special characters and end with either the `LF` or `CR+LF` Newline character. The system ignores any additional information in the entry.

Note the following:

- Deleting a source file from the file list also removes all associated SHA-256 hashes from the file list.
- You cannot upload multiple files to a file list if the successful source file upload results in the file list containing more than 10000 distinct SHA-256 values.
- The system truncates descriptions exceeding 256 characters to the first 256 characters on upload. If the description contains commas, you must use an escape character (`\,`). If no description is included, the source file name is used instead.
- If a file list contains a SHA-256 value, and you upload a source file containing that value, the newly uploaded value does not modify the existing SHA-256 value. When viewing captured files, file events, or malware events related to the SHA-256 value, any threat name or description is derived from the individual SHA-256 value.
- The system does not upload invalid SHA-256 values in a source file.
- If multiple uploaded source files contain an entry for the same SHA-256 value, the system uses the most recent value.
- If a source file contains multiple entries for the same SHA-256 value, the system uses the last one.
- You cannot directly edit a source file within the object manager. To make changes, you must first modify your source file directly, delete the copy on the system, then upload the modified source file. See Downloading a Source File from a File List, page 2-31 for more information.

To upload a source file to a file list:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Click **File List**.

Step 3 Click the edit icon () next to the file list where you want to add values from a source file.

Step 4 Choose `List of SHAs` from the **Add by** field.

Step 5 Optionally, enter a description of the source file in the **Description** field.

If you do not enter a description, the system uses the file name.

Step 6 Click **Browse** to browse to the source file, then click **Upload and Add List** to add the list.

The source file is added to the file list. The SHA-256 column lists how many SHA-256 values the file contains.

Step 7 Click **Store ASA FirePOWER Changes**.

Step 8 If an active policy references your object, deploy configuration changes; see Deploying Configuration Changes, page 4-11.

After configuration deployment, the system no longer performs malware cloud lookups on files in the file list.

Uploading an Individual File to a File List

License: Malware

If you have a copy of the file you want to add to a file list, you can upload the file to the system for analysis; the system calculates the file's SHA-256 value and adds the file to the list. The system does not enforce a limit on the size of files for SHA-256 calculation.

To add a file by having the system calculate its SHA-256 value:

Step 1 On the object manager's File List page, click the edit icon ( ) next to the clean list or custom detection list where you want to add a file.

Step 2 Choose **Calculate SHA** from the **Add by** field.

Step 3 Optionally, enter a description of the file in the **Description** field.

If you do not enter a description, the file name is used for the description on upload.

Step 4 Click **Browse** to browse to the source file, then click **Calculate and Add SHA** to add the list.

Step 5 Click **Store ASA FirePOWER Changes**.

Step 6 If an active policy references your object, deploy configuration changes; see Deploying Configuration Changes, page 4-11.

After configuration deployment, the system no longer performs malware cloud lookups on files in the file list.

Adding a SHA-256 Value to the File List

License: Malware

You can submit a file's SHA-256 value to add it to a file list. You cannot add duplicate SHA-256 values.

To add a file by manually entering the file's SHA-256 value:

Step 1 On the object manager's File List page, click the edit icon ( ) next to the clean list or custom detection list where you want to add a file.

Step 2 Choose `Enter SHA Value` from the **Add by** field.

Step 3 Enter a description of the source file in the **Description** field.

Step 4 Enter or paste the file's entire **SHA-256** value. The system does not support matching partial values.

Step 5 Click **Add** to add the file.

Step 6 Click **Store ASA FirePOWER Changes**.

Step 7 If an active policy references your object, deploy configuration changes; see Deploying Configuration Changes, page 4-11.

After configuration deployment, the system no longer performs malware cloud lookups on files in the file list.

Modifying Files on a File List

License: Malware

You can edit or delete individual SHA-256 values on a file list. Note that you cannot directly edit a source file within the object manager. To make changes, you must first modify your source file directly, delete the copy on the system, then upload the modified source file. See Downloading a Source File from a File List, page 2-31 for more information. To edit a file on a file list:

Step 1 On the object manager's File List page, click the edit icon ( ) next to the clean list or custom detection list where you want to modify a file.

Step 2 Next to the SHA-256 value you want to edit, click the edit icon ( ).

Tip You can also delete files from the list. Next to the file you want to remove, click the delete icon ( ).

Step 3 Update the **SHA-256** value or **Description**.

Step 4 Click **Save**.

Step 5 Click **Store ASA FirePOWER Changes**.

Step 6 If an active policy references your object, deploy configuration changes; see Deploying Configuration Changes, page 4-11.

After configuration deployment, the system no longer performs malware cloud lookups on files in the file list.

Downloading a Source File from a File List

License: Malware

You can view, download, or delete existing source file entries on a file list. Note that you cannot edit a source file once uploaded. You must first delete the source file from the file list, then upload an updated file. For more information on uploading a source file, see Uploading Multiple SHA-256 Values to a File List, page 2-28.

The number of entries associated with a source file refers to the number of distinct SHA-256 values. If you delete a source file from a file list, the total number of SHA-256 entries the file list contains decreases by the number of valid entries in the source file.

To download a source file:

Step 1 On the object manager's File List page, click the edit icon ( ) next to the clean list or custom detection list where you want to download a source file.

Step 2 Next to the source file you want to download, click the view icon ( ).

Step 3 Click **Download SHA List** and follow the prompts to save the source file.

Step 4 Click **Close**.

Working with Security Zones

License: Any

Supported Devices: Any

A *security zone* is a grouping of one or more ASA interfaces that you can use to manage and classify traffic flow in various policies and configurations. You can configure multiple zones on a single device. This allows you to divide the network into segments where the system can apply various policies. You must assign at least one interface to a security zone to match traffic against that security zone, and each interface can belong to only one zone.

In addition to using security zones to group interfaces, you can use zones in access control policies. For example, you could write an access control rule that applies only to a specific source or destination zone.

The Security Zones page of the object manager lists the zones configured on your ASA FirePOWER module.

You cannot delete a security zone that is in use. After you add or remove interfaces from a zone, if an active policy references your object, you must deploy the configuration to see your changes take effect; see Deploying Configuration Changes, page 4-11.

To create a security zone:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Choose **Security Zones**.

Step 3 Click **Add Security Zone**.

Step 4 Enter a **Name** for the zone. You can use any printable standard ASCII characters except curly braces ({}) and pound signs (#).

Step 5 Choose an interface **Type** for the zone.

After you create a security zone, you cannot change its type.

Step 6 Choose one or more interfaces.

Use the Shift and Ctrl keys to choose multiple objects. If you have not yet configured interfaces, you can create an empty zone and add interfaces to it later; skip to step 9.

Step 7 Click **Add**.

Step 8 Repeat steps 6 through 8 to add interfaces on other devices to the zone.

Step 9 Click **Store ASA FirePOWER Changes**.

Working with Cipher Suite Lists

License: Any

A cipher suite list is an object comprised of several cipher suites. Each pre-defined cipher suite value represents a cipher suite used to negotiate an SSL- or TLS-encrypted session. You can use cipher suites and cipher suite lists in SSL rules to control encrypted traffic based on whether the client and server negotiated the SSL session using that cipher suite. If you add a cipher suite list to an SSL rule, SSL sessions negotiated with any of the cipher suites in the list match the rule.

Note Although you can use cipher suites in the ASDM interface in the same places as cipher suite lists, you cannot add, modify, or delete cipher suites.

You cannot delete a cipher suite list that is in use. Additionally, after you edit a cipher suite list, if an active policy references your object, you must redeploy the configuration for your changes to take effect; see Deploying Configuration Changes, page 4-11.

To create a cipher suite list:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Choose **Cipher Suite List**.

Step 3 Click **Add Cipher Suites**.

Step 4 Enter a **Name** for the cipher suite list. You can use any printable standard ASCII characters except a pipe (|) or curly braces ({}).

Step 5 Choose one or more cipher suites and click **Add**.

- Use Shift and Ctrl to choose multiple cipher suites, or right-click and **Select All**.
- Use the filter field ( ) to search for existing cipher suites to include, which updates as you type to display matching items. Click the reload icon ( ) above the search field or click the clear icon ( ) in the search field to clear the search string.

Step 6 Click **Store ASA FirePOWER Changes**.

Working with Distinguished Name Objects

License: Any

Each distinguished name object represents the distinguished name listed for a public key certificate's subject or issuer. You can use distinguished name objects and groups (see Grouping Objects, page 2-2) in SSL rules to control encrypted traffic based on whether the client and server negotiated the SSL session using a server certificate with the distinguished name as subject or issuer.

Your distinguished name object can contain the common name attribute (**CN**). If you add a common name without "CN=" then the system prepends "CN=" before saving the object.

You can also add a distinguished name with one of each attribute listed in the following table, separated by commas.

Table 2-7 Distinguished Name Attributes

Attribute	Description	Allowed Values
C	Country Code	two alphabetic characters
CN	Common Name	up to 64 alphanumeric, backslash (/), hyphen (-), quotation ("), or asterisk (*) characters, or spaces
O	Organization	
OU	Organizational Unit	

You can define one or more asterisks (*) as wild cards in an attribute. In a common name attribute, you can define one or more asterisks per domain name label. Wild cards match only within that label, though you can define multiple labels with wild cards. See the following table for examples.

Table 2-8 Common Name Attribute Wild Card Examples

Attribute	Matches	Does Not Match
CN=”*ample.com”	example.com	mail.example.com example.text.com ampleexam.com
CN=”exam*.com”	example.com	mail.example.com example.text.com ampleexam.com
CN=”*xamp*.com”	example.com	mail.example.com example.text.com ampleexam.com
CN=”*.example.com”	mail.example.com	example.com example.text.com ampleexam.com
CN=”*.com”	example.com ampleexam.com	mail.example.com example.text.com
CN=”*.*.com”	mail.example.com example.text.com	example.com ampleexam.com

You cannot delete a distinguished name object that is in use. Additionally, after you edit a distinguished name object, if an active policy references your object, you must deploy the configuration for your changes to take effect; see Deploying Configuration Changes, page 4-11.

To create a distinguished name object:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **Distinguished Name**, choose **Individual Objects**.

Step 3 Click **Add Distinguished Name**.

Step 4 Enter a **Name** for the distinguished name object. You can use any printable standard ASCII characters except a pipe (|) or curly braces ({}).

Step 5 In the **DN** field, enter a value for the distinguished name or common name. You have the following options:

- If you add a distinguished name, you can include one of each attribute listed in Table 2-7 on page 2-33 separated by commas.
- If you add a common name, you can include multiple labels and wild cards.

Step 6 Click **Store ASA FirePOWER Changes**.

Working with PKI Objects

License: Any

PKI objects represent the public key certificates and paired private keys required to support your SSL inspection deployment. Internal and trusted CA objects consist of certificate authority (CA) certificates; internal CA objects also contain the private key paired with the certificate. Internal and external certificate objects consist of server certificates; internal certificate objects also contain the private key paired with the certificate. Using these objects in SSL rules, you can decrypt:

- outgoing traffic by re-signing the server certificate with an internal CA object
- incoming traffic using the known private key in an internal certificate object

You can also create SSL rules and match traffic encrypted with:

- the certificate in an external certificate object
- a certificate either signed by the CA in a trusted CA object, or within the CA's chain of trust

You can manually input certificate and key information, upload a file containing that information, or in some cases, generate a new CA certificate and private key.

When you view a list of PKI objects in the object manager, the system displays the certificate's Subject distinguished name as the object value. Hover your pointer over the value to view the full certificate Subject distinguished name. To view other certificate details, edit the PKI object.

Note The ASA FirePOWER module encrypts all private keys stored in internal CA objects and internal certificate objects with a randomly generated key before saving them. If you upload private keys that are password protected, the appliance decrypts the key using the user-supplied password, then reencrypts it with the randomly generated key before saving it.

For more information, see the following sections:

- Working with Internal Certificate Authority Objects, page 2-35
- Working with Trusted Certificate Authority Objects, page 2-39
- Working with External Certificate Objects, page 2-41
- Working with Internal Certificate Objects, page 2-41

Working with Internal Certificate Authority Objects

License: Any

Each internal certificate authority (CA) object you configure represents the CA public key certificate of a CA your organization controls. The object consists of the object name, CA certificate, and paired private key. You can use internal CA objects and groups (see Grouping Objects, page 2-2) in SSL rules to decrypt outgoing encrypted traffic by re-signing the server certificate with the internal CA.

Note If you reference an internal CA object in a **Decrypt - Resign** SSL rule and the rule matches an encrypted session, the user's browser may warn that the certificate is not trusted while negotiating the SSL handshake. To avoid this, add the internal CA object certificate to either the client or domain list of trusted root certificates.

You can create an internal CA object in the following ways:

- import an existing RSA-based or elliptic curve-based CA certificate and private key
- generate a new self-signed RSA-based CA certificate and private key
- generate an unsigned RSA-based CA certificate and private key. You must submit a certificate signing request (CSR) to another CA to sign the certificate before using the internal CA object.

After you create an internal CA object containing a signed certificate, you can download the CA certificate and private key. The system encrypts downloaded certificates and private keys with a user-provided password.

Whether system-generated or user-created, you can modify the internal CA object name, but cannot modify other object properties.

You cannot delete an internal CA object that is in use. Additionally, after you edit an internal CA object, if an active policy references your object, you must deploy the configuration for your changes to take effect; see Deploying Configuration Changes, page 4-11.

For more information, see the following sections:

- Importing a CA Certificate and Private Key, page 2-36
- Generating a New CA Certificate and Private Key, page 2-37
- Obtaining and Uploading a New Signed Certificate, page 2-37
- Downloading a CA Certificate and Private Key, page 2-38

Importing a CA Certificate and Private Key

License: Any

You can configure an internal CA object by importing an X.509 v3 CA certificate and private key. You can upload files encoded in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the private key file is password-protected, you can supply the decryption password. If the certificate and key are encoded in the PEM format, you can also copy and paste the information.

You can upload only files that contain proper certificate or key information, and that are paired with each other. The system validates the pair before saving the object.

Note If you configure a rule with the **Decrypt - Resign** action, the rule matches traffic based on the referenced internal CA certificate's encryption algorithm type, in addition to any configured rule conditions. You must upload an elliptic curve-based CA certificate to decrypt outgoing traffic encrypted with an elliptic curve-based algorithm, for example. For more information, see Decrypt Actions: Decrypting Traffic for Further Inspection, page 16-10.

To import an internal CA certificate and private key:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **PKI**, choose **Internal CAs**.

Step 3 Click **Import CA**.

Step 4 Enter a **Name** for the internal CA object. You can use any printable standard ASCII characters except a pipe (|) or curly braces ({}).

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 Above the **Key** field, click **Browse** to upload a DER or PEM-encoded paired private key file.

Step 7 If the uploaded file is password-protected, check the **Encrypted, and the password is:** check box and enter the password.

Step 8 Click **Store ASA FirePOWER Changes**.

The internal CA object is added.

Generating a New CA Certificate and Private Key

License: Any

You can configure an internal CA object by providing identification information to generate a self-signed RSA-based CA certificate and private key. The following table describes the identification information you provide to generate the certificate.

Table 2-9 Generated Internal CA Attributes

Field	Allowed Values	Required
Country Name (two-letter code)	two alphabetic characters	yes
State or Province	up to 64 alphanumeric, backslash (/), hyphen (-), quotation ("), asterisk (*), period (.), or space characters	no
Locality or City		
Organization		
Organizational Unit		
Common Name		

The generated CA certificate is valid for ten years. The Valid From date is a week before generation.

To generate a self-signed CA certificate:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **PKI**, choose **Internal CAs**.

Step 3 Click **Generate CA**.

Step 4 Enter a **Name** for the internal CA object. You can use any printable standard ASCII characters except a pipe (|) or curly braces ({}).

Step 5 Enter the identification attributes, as described in Table 2-9 on page 2-37.

Step 6 Click **Generate self-signed CA**.

Obtaining and Uploading a New Signed Certificate

License: Any

You can configure an internal CA object by obtaining a signed certificate from a CA. This involves two steps:

- Provide identification information to configure the internal CA object. This generates an unsigned certificate and paired private key, and creates a certificate signing request (CSR) to a CA you specify.
- After the CA issues the signed certificate, upload it to the internal CA object, replacing the unsigned certificate.

You can only reference an internal CA object in an SSL rule if it contains a signed certificate.

To create an unsigned CA certificate and CSR:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **PKI**, choose **Internal CAs**.

Step 3 Click **Generate CA**.

Step 4 Enter a **Name** for the internal CA object. You can use any printable standard ASCII characters except a pipe (|) or curly braces ({}).

Step 5 Enter the identification attributes, as described in Table 2-9 on page 2-37.

Step 6 Click **Generate CSR**.

Step 7 Copy the CSR to submit to a CA.

Step 8 Click **Store ASA FirePOWER Changes**.

Note that before you can use the CA, you must upload a signed certificate issued by a CA.

To upload a signed certificate issued in response to a CSR:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **PKI**, choose **Internal CAs**.

Step 3 Click the edit icon ( ) next to the CA object containing the unsigned certificate awaiting the CSR.

Step 4 Click **Install Certificate**.

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 If the uploaded file is password protected, check the **Encrypted, and the password is:** check box and enter the password.

Step 7 Click **Store ASA FirePOWER Changes**.

The CA object contains a signed certificate, and can be referenced in SSL rules.

Downloading a CA Certificate and Private Key

License: Any

You can back up or transfer a CA certificate and paired private key by downloading a file containing the certificate and key information from an internal CA object.

Caution Always store downloaded key information in a secure location.

The system encrypts the private key stored in an internal CA object with a randomly generated key before saving it to disk. If you download a certificate and private key from an internal CA object, the system first decrypts the information before creating a file containing the certificate and private key information. You must then provide a password the system uses to encrypt the downloaded file.

Caution Private keys downloaded as part of a system backup are decrypted, then stored in the unencrypted backup file. For more information, see Creating Backup Files, page 48-1.

To download an internal CA certificate and private key:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **PKI**, choose **Internal CAs**.

Step 3 Click the edit icon ( ) next to the internal CA object whose certificate and private key you want to download.

Step 4 Click **Download**.

Step 5 Enter an encryption password in the **Password** and **Confirm Password** fields.

Step 6 Click **Store ASA FirePOWER Changes**.

The system prompts you to save the file.

Working with Trusted Certificate Authority Objects

License: Any

Each trusted certificate authority (CA) object you configure represents a CA public key certificate belonging to a trusted CA outside your organization. The object consists of the object name and CA public key certificate. You can use external CA objects and groups (see Grouping Objects, page 2-2) in the SSL policy to control traffic encrypted with a certificate signed either by the trusted CA, or any CA within the chain of trust.

After you create the trusted CA object, you can modify the name and add certificate revocation lists (CRL), but cannot modify other object properties. There is no limit on the number of CRLs you can add to an object. If you want to modify a CRL you have uploaded to an object, you must delete the object and recreate it.

You cannot delete a trusted CA object that is in use. Additionally, after you edit a trusted CA object, if an active policy references your object, you must deploy the configuration for your changes to take effect; see Deploying Configuration Changes, page 4-11.

For more information, see the following sections:

- Working with Geolocation Objects, page 2-42
- Adding a Certificate Revocation List to a Trusted CA Object, page 2-40

Adding a Trusted CA Object

License: Any

You can configure an external CA object by uploading an X.509 v3 CA certificate. You can upload a file encoded in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the file is password-protected, you must supply the decryption password. If the certificate is encoded in the PEM format, you can also copy and paste the information.

You can upload a CA certificate only if the file contains proper certificate information; the system validates the certificate before saving the object.

To import a trusted CA certificate:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **PKI**, choose **Trusted CAs**.

Step 3 Click **Add Trusted CAs**.

Step 4 Enter a **Name** for the trusted CA object. You can use any printable standard ASCII characters except a pipe (`|`) or curly braces (`{}`).

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 If the file is password-protected, check the **Encrypted, and the password is:** check box and enter the password.

Step 7 Click **Store ASA FirePOWER Changes**.

Adding a Certificate Revocation List to a Trusted CA Object

License: Any

You can upload CRLs to a trusted CA object. If you reference that trusted CA object in an SSL policy, you can control encrypted traffic based on whether the CA that issued the session encryption certificate subsequently revoked the certificate. You can upload files encoded in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

After you add the CRL, you can view the list of revoked certificates. If you want to modify a CRL you have uploaded to an object, you must delete the object and recreate it.

You can upload only files that contain a proper CRL. There is no limit to the number of CRLs you can add to a trusted CA object. However, you must save the object each time you upload a CRL, before adding another CRL.

To upload a CRL:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **PKI**, choose **Trusted CAs**.

Step 3 Click the edit icon ( ) next to a trusted CA object.

Step 4 Click **Add CRL** to upload a DER or PEM-encoded CRL file.

Step 5 Click **Store ASA FirePOWER Changes**.

Working with External Certificate Objects

License: Any

Each external certificate object you configure represents a server public key certificate that does not belong to your organization. The object consists of the object name and certificate. You can use external certificate objects and groups (see Grouping Objects, page 2-2) in SSL rules to control traffic encrypted with the server certificate. For example, you can upload a self-signed server certificate that you trust, but cannot verify with a trusted CA certificate.

You can configure an external certificate object by uploading an X.509 v3 server certificate. You can upload a file in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

You can upload only files that contains proper server certificate information; the system validates the file before saving the object. If the certificate is encoded in the PEM format, you can also copy and paste the information.

After you create the external certificate object, you can modify the name, but cannot modify other object properties.

You cannot delete an external certificate object that is in use. Additionally, after you edit an external certificate object, if an active policy references your object, you must deploy the configuration for your changes to take effect; see Deploying Configuration Changes, page 4-11.

To create an external certificate object:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **PKI**, choose **External Certs**.

Step 3 Click **Add External Cert**.

Step 4 Enter a **Name** for the external certificate object. You can use any printable standard ASCII characters except a pipe (|) or curly braces ({}).

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 server certificate file.

Step 6 Click **Store ASA FirePOWER Changes**.

Working with Internal Certificate Objects

License: Any

Each internal certificate object you configure represents a server public key certificate belonging to your organization. The object consists of the object name, public key certificate, and paired private key. You can use internal certificate objects and groups (see Grouping Objects, page 2-2) in SSL rules to decrypt traffic incoming to one of your organization's servers using the known private key.

You can configure an internal certificate object by uploading an X.509 v3 RSA-based or elliptic curve-based server certificate and paired private key. You can upload a file in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the file is password-protected, you must supply the decryption password. If the certificate and key are encoded in the PEM format, you can also copy and paste the information.

You can upload only files that contain proper certificate or key information, and that are paired with each other. The system validates the pair before saving the object.

After you create the internal certificate object, you can modify the name, but cannot modify other object properties.

You cannot delete an internal certificate object that is in use. Additionally, after you edit an internal certificate object, if an active policy references your object, you must deploy the configuration for your changes to take effect; see Deploying Configuration Changes, page 4-11.

To create an internal certificate object:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Under **PKI**, choose **Internal Certs**.

Step 3 Click **Add Internal Cert**.

Step 4 Enter a **Name** for the internal certificate object. You can use any printable standard ASCII characters except a pipe (`|`) or curly braces (`{}`).

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 server certificate file.

Step 6 Above the **Key** field, or click **Browse** to upload a DER or PEM-encoded paired private key file.

Step 7 If the uploaded private key file is password-protected, check the **Encrypted, and the password is:** check box and enter the password.

Step 8 Click **Store ASA FirePOWER Changes**.

Working with Geolocation Objects

License: Any

Each geolocation object you configure represents one or more countries or continents that the system has identified as the source or destination of traffic on your monitored network. You can use geolocation objects in access control policies or SSL policies. For example, you could write an access control rule that blocks traffic to or from certain countries. For information on filtering traffic by geographical location, see Controlling Traffic by Network or Geographical Location, page 7-3.

To ensure that you are using up-to-date information to filter your network traffic, Cisco strongly recommends that you regularly update your Geolocation Database (GeoDB). For information on downloading and installing GeoDB updates, see Updating the Geolocation Database, page 46-19.

You cannot delete a geolocation object that is in use. Additionally, after you edit a geolocation object used in an access control policy or SSL policy, you must redeploy policies for your changes to take effect.

To create a geolocation object:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Choose **Geolocation**.

Step 3 Click **Add Geolocation**.

Step 4 Enter a **Name** for the geolocation object. You can use any printable standard ASCII characters except curly braces ({}).

Step 5 Check the check boxes for the countries and continents you want to include in your geolocation object.

Selecting a continent selects all countries within that continent, as well as any countries that GeoDB updates may add under that continent in the future. Deselecting any country under a continent deselects the continent. You can select any combination of countries and continents.

Step 6 Click **Store ASA FirePOWER Changes**.

Working with Security Group Tag Objects

License: Any

A Security Group Tag (SGT) object specifies a single SGT value, which you can use as a custom SGT condition in access control rules. You cannot group SGT objects.

If you configure ISE/ISE-PIC as an identity source, the system automatically disables the Security Group Tag option in the Object Manager. You cannot add new SGT objects, edit existing SGT objects, or use SGT objects as rule conditions unless you disable the ISE/ISE-PIC connection. For more information on the difference between custom SGTs and ISE SGTs, see ISE SGT v. Custom SGT Rule Conditions, page 10-1.

If you edit or delete an SGT object, if an active policy references the object, you must redeploy your configuration for the changes to take effect; see Deploying Configuration Changes, page 4-11.

To create an SGT object:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Object Management**.

Step 2 Choose **Security Group Tag**.

Step 3 Click **Add Security Group Tag**.

Step 4 Enter a **Name**.

Step 5 Optionally, enter a **Description**.

Step 6 In the **Tag** field, enter a single SGT.

Step 7 Click **Store ASA FirePOWER Changes**.



Managing Device Configuration

The Device Management page allows you to manage the device and interface configurations for the ASA FirePOWER module.

Caution If you configure the ASA in a failover pair, the ASA FirePOWER configuration does not automatically synchronize with the ASA FirePOWER module on the secondary device. You must manually export the ASA FirePOWER configuration from the primary and import it into the secondary every time you make a change.

For more information, see the following sections:

- Editing Device Configuration, page 3-1
- Managing ASA FirePOWER Module Interfaces, page 3-4
- Applying Changes to Device Configuration, page 3-4
- Configuring Remote Management, page 3-5
- Configuring eStreamer on the eStreamer Server, page 3-7

Editing Device Configuration

The Device tab of the Device Management page displays detailed device configuration and information, as it applies to the ASA FirePOWER module. It also allows you to make changes to some parts of device configuration, such as changing the displayed module name and modifying management settings.

See the following sections for more information:

- Editing General Device Configuration, page 3-1
- Viewing Device System Settings, page 3-2
- Understanding Advanced Device Settings, page 3-2

Editing General Device Configuration

License: Any

The General section of the Device tab shows the module name, which you can change. Here, you can also specify whether or not the device can transfer packets to the ASA FirePOWER module.

To edit general device configuration:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Device Management > Device**.

The Device page appears.

Step 2 Next to the **General** section, click the edit icon ( ).

The General pop-up window appears.

Step 3 In the **Name** field, type a new assigned name for the module. You may enter alphanumeric characters and special characters, with the exception of the following characters, which are invalid: +, (,), {, }, #, &, \, <, >, ?, ‘, and “.

Step 4 Select the **Transfer Packets** check box to allow packet data to be stored in the ASA FirePOWER module with events. Clear the check box to prevent the device from sending packet data with the events.

Step 5 Click **Save**.

The changes are saved. Note that your changes do not take effect until you apply the device configuration; see Applying Changes to Device Configuration, page 3-4 for more information.

Viewing Device System Settings

License: Any

The System section of the Device tab displays a read-only table of system information, as described in the following table.

Table 3-1 System Section Table Fields

Field	Description
Model	The model name and number for the device.
Serial	The serial number of the chassis of the device.
Time	The current system time of the device.
Version	The version of the software currently installed on the ASA FirePOWER module.
Policy	A link to the system policy currently applied to the ASA FirePOWER module.

Understanding Advanced Device Settings

The Advanced section of the Device tab displays advanced configuration settings, as described in the following table.

Table 3-2 Advanced Section Table Fields

Field	Description
Application Bypass	The state of Automatic Application Bypass on the module.
Bypass Threshold	The Automatic Application Bypass threshold, in milliseconds.

You can use the Advanced section to edit any of these settings. See the following sections for more information:

- Automatic Application Bypass, page 3-3
- Editing Advanced Device Settings, page 3-3

Automatic Application Bypass

License: Any

The Automatic Application Bypass (AAB) feature limits the time allowed to process packets through an interface and allows packets to bypass detection if the time is exceeded. The feature functions with any deployment; however, it is most valuable in inline deployments.

You balance packet processing delays with your network's tolerance for packet latency. When a malfunction within Snort or a device misconfiguration causes traffic processing time to exceed a specified threshold, AAB causes Snort to restart within ten minutes of the failure, and generates troubleshoot data that can be analyzed to investigate the cause of the excessive processing time.

You can change the bypass threshold if the option is selected. The default setting is 3000 milliseconds (ms). The valid range is from 250 ms to 60,000 ms.

Note AAB is activated only when an excessive amount of time is spent processing a single packet. If AAB engages, the system kills all Snort processes.

For more information about enabling Automatic Application Bypass and setting the bypass threshold, see Editing Advanced Device Settings, page 3-3.

Editing Advanced Device Settings

You can use the Advanced section of the Devices tab to modify the Automatic Application Bypass.

To modify advanced device settings:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Device Management > Device**.

The Device page appears.

Step 2 Next to the **Advanced** section, click the edit icon ( ).

The Advanced pop-up window appears.

Step 3 Optionally, select **Automatic Application Bypass** if your network is sensitive to latency. Automatic Application Bypass is most useful in inline deployments. For more information, see Automatic Application Bypass, page 3-3.

Step 4 When you select the Automatic Application Bypass option, you can type a **Bypass Threshold** in milliseconds (ms). The default setting is 3000 ms and the valid range is from 250 ms to 60,000 ms.

Step 5 Click **Save**.

Your changes are saved. Note that your changes do not take effect until you apply the device configuration; see Applying Changes to Device Configuration, page 3-4 for more information.

Managing ASA FirePOWER Module Interfaces

License: Control, Protection

When editing an ASA FirePOWER interface, you can configure only the interface's security zone from the ASA FirePOWER module. See Working with Security Zones, page 2-32 for more information.

You configure interfaces using ASDM and CLI.

To edit an ASA FirePOWER Interface:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Device Management > Interfaces**.

The Interfaces page appears.

Step 2 Next to the interface you want to edit, click the edit icon ( ).

The Edit Interface pop-up window appears.

Step 3 From the **Security Zone** drop-down list, select an existing security zone or select **New** to add a new security zone.

Step 4 Click **Store ASA FirePOWER Changes**.

The security zone is configured. Note that your changes do not take effect until you apply the device configuration; see Applying Changes to Device Configuration, page 3-4 for more information.

Applying Changes to Device Configuration

License: Any

After you make changes to the ASA FirePOWER configuration of a device, you must apply the changes before they take effect throughout the module. Note that the device must have unapplied changes or this option remains disabled.

Note that if you edit interfaces and reapply a device policy, Snort restarts for all interface instances on the device, not just those that you edited.

To apply changes to the device:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Device Management > Device** or **Configuration > ASA FirePOWER Configuration > Device Management > Interfaces**.

The Device Management page appears.

Step 2 Click **Apply ASA FirePOWER Changes**.

Step 3 When prompted, click **Apply**.

The device changes are applied.

Tip Optionally, from the Apply Device Changes dialog box, click **View Changes**. The Device Management Revision Comparison Report page appears in a new window. For more information, see Using the Device Management Revision Comparison Report, page 3-5.

Step 4 Click **OK**.

You are returned to the Device Management page.

Using the Device Management Revision Comparison Report

License: Any

A device management comparison report allows you to view the changes you have made to an appliance before you apply them. The report displays all differences between the current appliance configuration and the proposed appliance configuration. This gives you an opportunity to discover any potential configuration errors.

To compare appliance changes before applying them:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Device Management > Device** or **Configuration > ASA FirePOWER Configuration > Device Management > Interfaces**.

The Device Management page appears.

Step 2 Click **Apply Changes**.

The Apply Device Changes pop-up window appears. Note that the appliance must have unapplied changes or the Apply Changes button remains disabled.

Step 3 Click **View Changes**.

The Device Management Revision Comparison Report page appears in a new window.

Step 4 Click **Previous** and **Next** to scroll through the differences between the current appliance configuration and the proposed appliance configuration.

Step 5 Optionally, click **Comparison Report** to produce a PDF version of the report.

Configuring Remote Management

License: Any

Before you can manage one Firepower system appliance with another, you must set up a two-way, SSL-encrypted communication channel between the two appliances. The appliances use the channel to share configuration and event information. High availability peers also use the channel, which is by default on port 8305/tcp.

You must configure remote management on the appliance that will be managed, that is, on the device that you want to manage with a Firepower Management Center. After you configure remote management, you can use the managing appliance's web interface to add the managed appliance to your deployment.

Note After you establish remote management and register the Cisco ASA with FirePOWER Services to a Firepower Management Center, you **must** manage the ASA FirePOWER module from the Firepower Management Center instead of ASDM. You cannot remotely manage the Cisco ASA with FirePOWER Services with the ASDM console once the appliance is registered to a Firepower Management Center.

To enable communications between two appliances, you must provide a way for the appliances to recognize each other. There are three criteria the Firepower system uses when allowing communications:

- the hostname or IP address of the appliance with which you are trying to establish communication

  In NAT environments, even if the other appliance does not have a routable address, you must provide a hostname or an IP address either when you are configuring remote management, or when you are adding the managed appliance.
- a self-generated alphanumeric registration key up to 37 characters in length that identifies the connection
- an optional unique alphanumeric NAT ID that can help the Firepower system establish communications in a NAT environment

  The NAT ID **must** be unique among all NAT IDs used to register managed appliances.

When you register a managed device to a Firepower Management Center, the access control policy you select applies to the device. However, if you do not enable licenses for the device required by features used in the access control policy you select, the access control policy apply fails.

To configure remote management of the local appliance:

Access: Admin

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > Configuration > Registration**.

The Remote Management page appears.

Step 2 Click **Add Manager**.

The Add Remote Management page appears.

Step 3 In the **Management Host** field, type the IP address or the hostname of the appliance that you want to use to manage this appliance.

The hostname is the fully qualified domain name or the name that resolves through the local DNS to a valid IP address.

In a NAT environment, you do not need to specify an IP address or hostname here if you plan to specify it when you add the managed appliance. In this case, the Firepower system uses the NAT ID you will provide later to identify the remote manager on the managed ASA FirePOWER module interface.

Caution Use a hostname rather than an IP address if your network uses DHCP to assign IP addresses.

Step 4 In the **Registration Key** field, type the registration key that you want to use to set up communications between appliances.

Step 5 For NAT environments, in the **Unique NAT ID** field, type a **unique** alphanumeric NAT ID that you want to use to set up communications between appliances.

Step 6 Click **Save**.

After the appliances confirm that they can communicate with each other, the Pending Registration status appears.

Step 7 Use the managing appliance's web user interface to add this appliance to your deployment.

Note When enabling remote management of a device, in some high availability deployments that use NAT, you may also need to add the secondary Firepower Management Center as a manager. For more information, contact Support.

Editing Remote Management

License: Any

Use the following procedure to edit the hostname or IP address of the managing appliance. You can also change the display name of the managing appliance, which is a name only used within the context of the Firepower system. Although you can use the hostname as the display name of the appliance, entering a different display name does not change the hostname.

To edit remote management:

Access: Admin

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > Configuration > Registration**.

The Remote Management page appears.

Step 2 Click the edit icon ( ) next to the manager for which you want to edit remote management settings.

The Edit Remote Management page appears.

Step 3 In the **Name** field, change the display name of the managing appliance.

Step 4 In the **Host** field, change the IP address or the hostname of the managing appliance.

The hostname is the fully qualified domain name or the name that resolves through the local DNS to a valid IP address.

Step 5 Click **Save**.

Your changes are saved.

Configuring eStreamer on the eStreamer Server

License: FireSIGHT + Protection

Before the appliance you want to use as an eStreamer server can begin streaming eStreamer events to an external client, you must configure the eStreamer server to send events to clients, provide information about the client, and generate a set of authentication credentials to use when establishing communication.

Configuring eStreamer Event Types

You can control which types of events the eStreamer server is able to transmit to clients that request them.

Available event types on either a managed device or a Firepower Management Center are:

- Intrusion events
- Intrusion event packet data
- Intrusion event extra data

To configure the types of events transmitted by eStreamer:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > Configuration > Registration**.

The Registration page appears.

Step 2 Select the **eStreamer** tab.

The eStreamer page appears.

Step 3 Under **eStreamer Event Configuration**, select the check boxes next to the types of events you want eStreamer to forward to requesting clients.

You can select any or all of the following on a managed device or Firepower Management Center:

- **Intrusion Events** to transmit intrusion events.
- **Intrusion Event Packet Data** to transmit packets associated with intrusion events.
- **Intrusion Event Extra Data** to transmit additional data associated with an intrusion event such as the originating IP addresses of a client connecting to a web server through an HTTP proxy or load balancer.

Note Note that this controls which events the eStreamer server can transmit. Your client must still specifically request the types of events you want it to receive in the request message it sends to the eStreamer server. For more information, see the *Firepower system eStreamer Integration Guide*.

Step 4 Click **Save**.

Your settings are saved and the events you selected will be forwarded to eStreamer clients when requested.

Adding Authentication for eStreamer Clients

Before eStreamer can send eStreamer events to a client, you must add the client to the eStreamer server's peers database from the eStreamer page. You must also copy the authentication certificate generated by the eStreamer server to the client.

To add an eStreamer client:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > Configuration > Registration**.

The Registration page appears.

Step 2 Select the **eStreamer** tab.

The eStreamer page appears.

Step 3 Click **Create Client**.

The Create Client page appears.

Step 4 In the **Hostname** field, enter the host name or IP address of the host running the eStreamer client.

Note If you use a host name, the eStreamer server **must** be able to resolve the host to an IP address. If you have not configured DNS resolution, you should configure it first or use an IP address.

Step 5 If you want to encrypt the certificate file, enter a password in the **Password** field.

Step 6 Click **Save**.

The eStreamer server now allows the host to access port 8302 on the eStreamer server and creates an authentication certificate to use during client-server authentication. The eStreamer page reappears, with the new client listed under **Hostname**.

Step 7 Click the download icon ( ) next to the client hostname to download the certificate file.

Step 8 Save the certificate file to the appropriate directory used by your client for SSL authentication.

The client can now connect to the eStreamer server. You do not need to restart the eStreamer service.

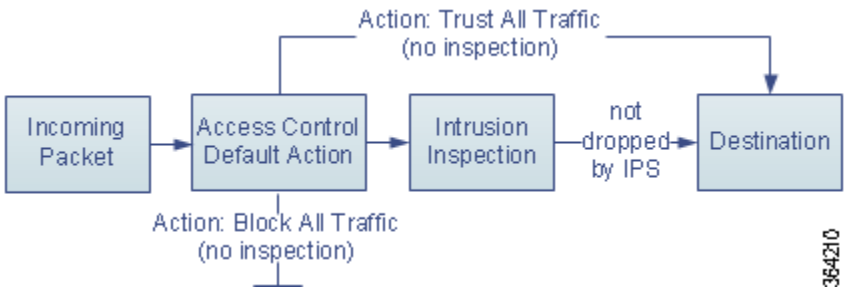
Tip To revoke access for a client, click the delete icon ( ) next to the host you want to remove. Note that you do not need to restart the eStreamer service; access is revoked immediately.



Getting Started with Access Control Policies

An *access control policy* determines how the system handles traffic on your network. Each ASA FirePOWER module can have one currently applied policy.

The simplest access control policy handles all traffic using its *default action*. You can set this default action to block or trust all traffic without further inspection, or to inspect traffic for intrusions.



Note that only ASA FirePOWER modules deployed inline can affect the flow of traffic. Applying an access control policy configured to block or alter traffic to passively deployed devices can have unexpected results. In some cases, the system prevents you from applying inline configurations to passively deployed ASA FirePOWER modules.

This chapter explains how to create and apply a simple access control policy. It also contains basic information on managing access control policies: editing, updating, comparing, and so on. For more information, see:

- Access Control License and Role Requirements, page 4-2
- Creating a Basic Access Control Policy, page 4-3
- Managing Access Control Policies, page 4-6
- Editing Access Control Policies, page 4-7
-
- Understanding Out-of-Date Policy Warnings, page 4-10
- Deploying Configuration Changes, page 4-11
- Troubleshooting Access Control Policies and Rules, page 4-12
- Generating a Report of Current Access Control Settings, page 4-15
- Comparing Access Control Policies, page 4-16

A more complex access control policy can blacklist traffic based on Security Intelligence data, as well as use *access control rules* to exert granular control over network traffic logging and handling. These rules can be simple or complex, matching and inspecting traffic using multiple criteria. Advanced access control policy options control decryption, preprocessing, performance, and other general preferences.

After you create a basic access control policy, see the following chapters for more information on tailoring it to your deployment:

- Blacklisting Using Security Intelligence IP Address Reputation, page 5-1 explains how to immediately blacklist (block) connections based on the latest reputation intelligence.
- Understanding Network Analysis and Intrusion Policies, page 18-1 explains how network analysis and intrusion policies preprocess and examine packets, as part of the system's intrusion detection and prevention feature.
- Tuning Traffic Flow Using Access Control Rules, page 6-1 explains how access control rules provide a granular method of handling network traffic across multiple ASA FirePOWER modules.
- Controlling Traffic Using Intrusion and File Policies, page 11-1 explains how intrusion and file policies provide the last line of defense before traffic is allowed to its destination, by detecting and optionally blocking intrusions, prohibited files, and malware.

Access Control License and Role Requirements

Although you can create access control policies regardless of the licenses on your ASA FirePOWER module, many features require that you enable the appropriate licenses before you apply the policy.

For more information, see License Requirements for Access Control, page 4-2.

License Requirements for Access Control

Although you can create access control policies regardless of the licenses on your ASA FirePOWER module, certain aspects of access control require that you enable specific licensed capabilities before you can apply the policy.

Warning icons and confirmation dialog boxes designate unsupported features for your deployment. For details, see Troubleshooting Access Control Policies and Rules, page 4-12.

The following table explains the license requirements to apply access control policies.

Table 4-1 License Requirements for Access Control

To apply an access control policy that...	License
performs access control based on zone, network, or port performs URL filtering using literal URLs and URL objects	Any
performs access control using geolocation data (source or destination country or continent)	Any
performs intrusion detection and prevention, file control, or Security Intelligence filtering	Protection
performs advanced malware protection, that is, network-based malware detection and blocking	Malware

Table 4-1 License Requirements for Access Control (continued)

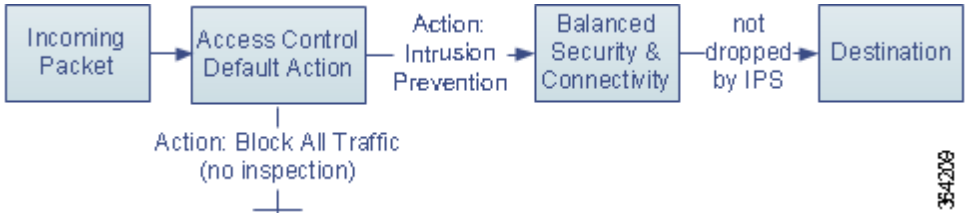
To apply an access control policy that...	License
performs user or application control	Control
performs URL filtering using category and reputation data	URL Filtering

Creating a Basic Access Control Policy

License: Any

When you create a new access control policy you must give it a unique name and specify a default action. At this point, the default action determines how the ASA FirePOWER module handles all unencrypted traffic; you will add other configurations that affect traffic flow later.

When you create a new policy, you can set the default action to block all traffic without further inspection, or to inspect traffic for intrusions, as shown in the following diagram.



Tip When you first create an access control policy, you cannot choose to trust traffic as the default action. If you want to trust all traffic by default, change the default action after you create the policy.

Use the Access Control Policy page (**Policies > Access Control**) to create new and manage existing access control policies.

Optionally, you can use and modify the initial system-provided policy named Default Trust All Traffic.

To create an access control policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Tip You can also copy an existing policy from this ASA FirePOWER module or import a policy from another ASA FirePOWER module. To copy a policy, click the copy icon ( ). To import a policy, see Importing and Exporting Configurations, page B-1.

Step 2 Click **New Policy**.

The New Access Control Policy pop-up window appears.

Step 3 Give the policy a unique **Name** and, optionally, a **Description**.

You can use all printable characters, including spaces and special characters, except for the pound sign (#), a semi-colon (;), or either brace ({}). The name must include at least one non-space character.

Step 4 Specify the initial **Default Action**:

- **Block all traffic** creates a policy with the **Access Control: Block All Traffic** default action.
- **Intrusion Prevention** creates a policy with the **Intrusion Prevention: Balanced Security and Connectivity** default action.

For guidance on choosing an initial default action, as well as how to change it later, see Setting Default Handling and Inspection for Network Traffic, page 4-4.

Step 5 Click **Store ASA FirePOWER Changes**.

The access control policy editor appears. For information on configuring your new policy, see Editing Access Control Policies, page 4-7. Note that you must apply the policy for it to take effect; see Deploying Configuration Changes, page 4-11.

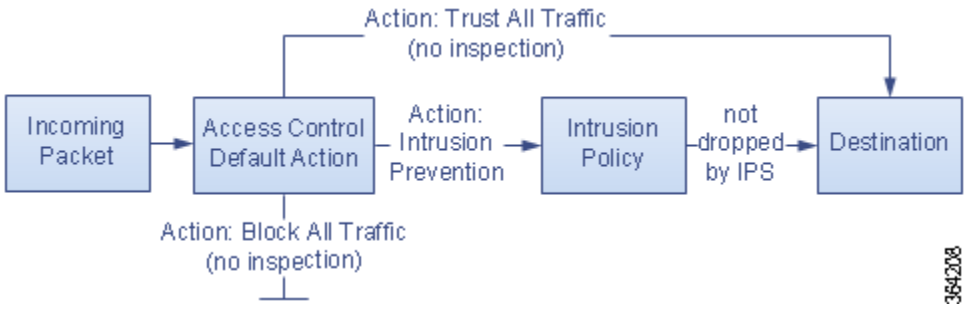
Setting Default Handling and Inspection for Network Traffic

License: Any

When you create an access control policy, you must select a default action. The default action for an access control policy determines how the system handles decrypted or unencrypted traffic that:

- is not blacklisted by Security Intelligence
- matches none of the rules in the policy (except Monitor rules, which match and log—but do not handle or inspect—traffic)

Therefore, when you apply an access control policy that does not contain any access control rules or Security Intelligence configurations, and that does not invoke an SSL policy to handle encrypted traffic, the default action determines how *all* traffic on your network is handled. You can block or trust all traffic without further inspection, or inspect traffic for intrusions. Your options are shown in the following diagram.

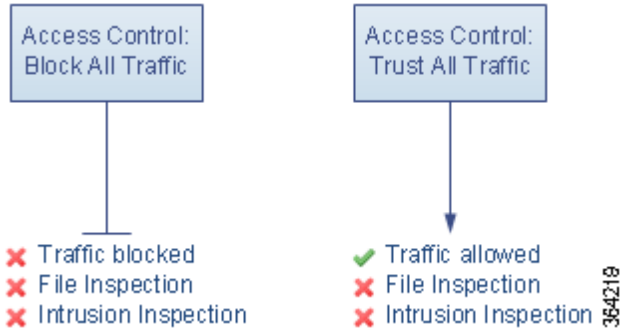


The following table describes how the different default actions handle traffic, and lists the types of inspection you can perform on traffic handled by each default action. Note that you **cannot** perform file or malware inspection on traffic handled by the default action. For more information, see Controlling Traffic Using Intrusion and File Policies, page 11-1.

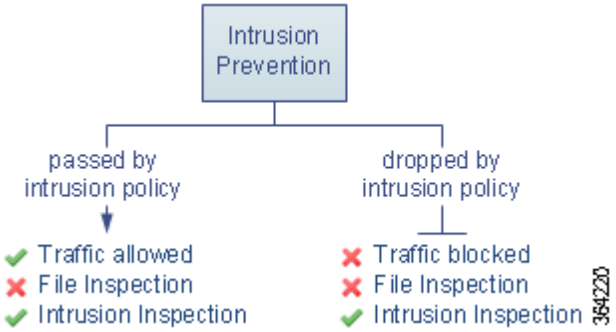
Table 4-2 Access Control Policy Default Actions

Default Action	Effect on Traffic	Inspection Type and Policy
Access Control: Block All Traffic	block without further inspection	none
Access Control: Trust All Traffic	trust (allow to its final destination without further inspection)	none
Intrusion Prevention	allow, as long as it is passed by the intrusion policy you specify (requires a Protection license)	intrusion, using the specified intrusion policy and associated variable set

The diagram below illustrates the **Block All Traffic** and **Trust All Traffic** default actions.



The diagram below illustrates the **Intrusion Prevention** default actions.



When you first create an access control policy, logging connections that are handled by the default action is disabled by default. If you select a default action that performs intrusion inspection, the system automatically associates the default intrusion variable set with the intrusion policy you select. You can change either of these options, as well as the default action itself, after you create the policy.

To change an access control policy's default action and related options:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to configure.

The access control policy editor appears.

Step 3 Select a **Default Action**.

- To block all traffic, select **Access Control: Block All Traffic**.
- To trust all traffic, select **Access Control: Trust All Traffic**.
- To inspect all traffic with an intrusion policy, select an intrusion policy, all of which begin with the label **Intrusion Prevention**. Keep in mind that an intrusion policy can block traffic.

Caution Do **not** use `Experimental Policy 1` unless instructed to do so by a Cisco representative. Cisco uses this policy for testing.

Step 4 If you selected an **Intrusion Prevention** default action, click the variables icon ( ) to change the variable set associated with the intrusion policy you selected.

In the pop-up window that appears, select a new variable set and click **OK**. You can also edit the selected variable set in a new window by clicking the edit icon ( ). If you do not change the variable set, the system uses a default set. For more information, see Working with Variable Sets, page 2-13.

Step 5 Click the logging icon ( ) to change logging options for connections handled by the default action.

You can log a matching connection at its beginning and end. Note that the system cannot log the end of blocked traffic. You can log connections to the ASA FirePOWER module event viewer, external system log (syslog) or SNMP trap server. For more information, see Logging Connections Handled by the Access Control Default Action, page 36-11.

Managing Access Control Policies

License: Any

On the Access Control Policy page (**Configuration > ASA FirePOWER Configuration > Policies > Access Control**) you can view your current custom access control policies, along with information on whether a policy is applied.

In addition to custom policies that you create, the system provides a custom policy Default Allow All Traffic that you can edit and use.

Options on the Access Control Policy page allow you to take the actions in the following table.

Table 4-3 Access Control Policy Management Actions

To...	You can...	See...
create a new access control policy	click **New Policy**.	Creating a Basic Access Control Policy, page 4-3
edit an existing access control policy	click the edit icon ( ).	Editing Access Control Policies, page 4-7
reapply an access control policy	click the apply icon ( ).	Deploying Configuration Changes, page 4-11

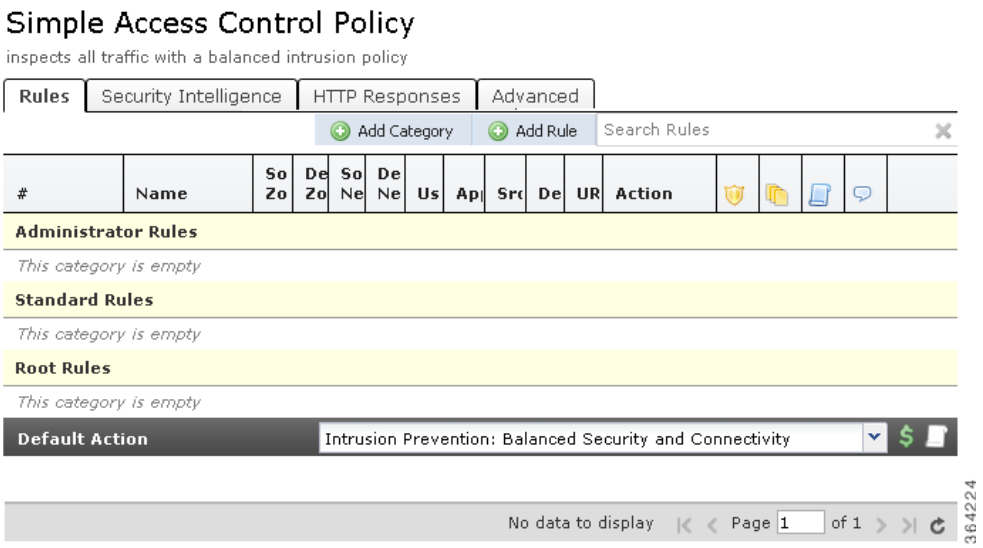
Table 4-3 Access Control Policy Management Actions (continued)

To...	You can...	See...
export an access control policy to import on another ASA FirePOWER module	click the export icon ( )	Exporting Configurations, page B-1
view a PDF report that lists the current configuration settings in an access control policy	click the report icon ( ).	Generating a Report of Current Access Control Settings, page 4-15
compare access control policies	click **Compare Policies**.	Comparing Access Control Policies, page 4-16
delete an access control policy	click the delete icon ( ), then confirm that you want to delete the policy. You cannot delete an applied access control policy or one that is currently applying.	

Editing Access Control Policies

License: Any

When you first create a new access control policy, the access control policy editor appears, focused on the Rules tab. The following graphic shows a newly created policy. Because a new policy does not yet have rules or other configurations, the default action handles *all* unencrypted traffic. In this case, the default action inspects traffic with the system-provided Balanced Security and Connectivity intrusion policy before allowing it to its final destination.



Use the access control policy editor to add and organize rules, and so on. The following list provides information on the policy configurations you can change.

Name and Description

To change the policy's name and description, click the appropriate field and type the new name or description.

Security Intelligence

Security Intelligence is a first line of defense against malicious Internet content. This feature allows you to immediately blacklist (block) connections based on the latest reputation intelligence. To ensure continual access to vital resources, you can override blacklists with custom whitelists. This traffic filtering takes place **before** any other policy-based inspection, analysis, or traffic handling, including rules and the default action. For more information, see Blacklisting Using Security Intelligence IP Address Reputation, page 5-1.

Rules

Rules provide a granular method of handling network traffic. Rules in an access control policy are numbered, starting at 1. The system matches traffic to access control rules in top-down order by ascending rule number.

In most cases, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic. These conditions include security zone, network or geographical location, port, application, requested URL, or user. Conditions can be simple or complex; their use often depends on certain licenses.

Use the Rules tab to add, categorize, enable, disable, filter, and otherwise manage rules. For more information, see Tuning Traffic Flow Using Access Control Rules, page 6-1.

Default Action

The default action determines how the system handles traffic that is not blacklisted by Security Intelligence and does not match any access control rules. Using the default action, you can block or trust all traffic without further inspection, or inspect traffic for intrusions. You can also enable or disable logging of connections handled by the default action.

For more information, see Setting Default Handling and Inspection for Network Traffic, page 4-4 and Logging Connections Based on Access Control Handling, page 36-9.

HTTP Responses

You can specify what the user sees in a browser when the system blocks that user's website request—either display a generic system-provided response page, or enter custom HTML. You can also display a page that warns users, but also allows them to click a button to continue or refresh the page to load the originally requested site. For more information, see Displaying a Custom Web Page for Blocked URLs, page 8-14.

Advanced Access Control Options

Advanced access control policy settings typically require little or no modification. The default settings are appropriate for most deployments. Advanced settings you can modify include:

- the number of characters you store in the ASA FirePOWER module database for each URL requested by your users; see Logging URLs Detected in Connections, page 36-13
- the length of time before you re-block a website after a user bypasses an initial block; see Setting the User Bypass Timeout for a Blocked Website, page 8-13
- network analysis and intrusion policy settings that allow you to tailor many preprocessing options to networks and zones, as well as set default intrusion inspection behavior; see Customizing Traffic Preprocessing, page 20-1
- advanced transport and network preprocessor settings that apply globally to all networks and zones where you apply the access control policy; see Configuring Advanced Transport/Network Settings, page 24-1

- adaptive profiles to improve reassembly of packet fragments and TCP streams in passive deployments, based on your network's host operating systems; see Tuning Preprocessing in Passive Deployments, page 25-1
- performance options for intrusion inspection, file control, and advanced malware protection; see Tuning Intrusion Prevention Performance, page 11-6 and Tuning File and Malware Inspection Performance and Storage, page 11-16

When you edit an access control policy, a message indicates that you have unsaved changes. To retain your changes, you must save the policy before exiting the policy editor. If you attempt to exit the policy editor without saving your changes, you are cautioned that you have unsaved changes; you can then discard your changes and exit the policy, or return to the policy editor.

To protect the privacy of your session, after sixty minutes of inactivity on the policy editor, changes to your policy are discarded and you are returned to the Access Control Policy page. After the first thirty minutes of inactivity, a message appears and updates periodically to provide the number of minutes remaining before changes are discarded. Any activity on the page cancels the timer.

To edit an access control policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to configure.

The access control policy editor appears.

Step 3 Edit your policy. Take any of the actions summarized above.

Step 4 Save or discard your configuration:

- To save your changes and continue editing, click **Store ASA FirePOWER Changes.**
- To save your changes and apply your policy, click **Apply ASA FirePOWER Changes**. See Deploying Configuration Changes, page 4-11.
- To discard your changes, click **Cancel** and, if prompted, click **OK**.

Associating Other Policies with Access Control

License: Any

Use an access control policy's advanced settings to associate one of each of the following subpolicies with the access control policy:

- SSL policy—Monitors, decrypts, blocks, or allows application layer protocol traffic encrypted with Secure Socket Layer (SSL) or Transport Layer Security (TLS) .
- Identity policy—Performs user authentication based on the realm and authentication method associated with the traffic.

Caution Associating an SSL or identity policy, or subsequently dissociating the policy by choosing None, restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on the model of the managed device and how it handles traffic.

To associate other policies with an access control policy:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

Step 2 Click the edit icon ( ) next to the access control policy you want to configure.

Step 3 Click the **Advanced** tab.

Step 4 Click the edit icon ( ) in the appropriate Policy Settings area.

Step 5 Choose a policy from the drop-down list.

If you choose a user-created policy, you can edit the policy by clicking the edit icon.

Step 6 Click **OK**.

Step 7 Save or discard your configuration:

- To save your changes and continue editing, click **Store ASA FirePOWER Changes.**
- To save your changes and apply your policy, click **Apply ASA FirePOWER Changes**. See Deploying Configuration Changes, page 4-11.
- To discard your changes, click **Cancel** and, if prompted, click **OK**.

Understanding Out-of-Date Policy Warnings

License: Any

On the Access Control Policy page (**Configuration > ASA FirePOWER Configuration > Policies > Access Control**), out-of-date policies are marked with red status text.

In almost every case, whenever you change an access control policy, you must reapply it for the change to take effect. If the access control policy invokes other policies or relies on other configurations, changing those also requires that you reapply the access control policy (or, for intrusion policy changes, you can reapply just the intrusion policy).

Configuration changes that require a policy reapply include:

- Modifying the access control policy itself: any changes to access control rules, the default action, Security Intelligence filtering, advanced options including NAP rules, and so on.
- Changing any of the intrusion and file policies that the access control policy invokes: the SSL policy, network analysis policies, intrusion policies, and file policies.
- Changing any reusable object or configuration used in the access control policy or the policies it invokes: network, port, URL, and geolocation objects; Security Intelligence lists and feeds; application filters or detectors; intrusion policy variable sets; file lists; decryption-related objects, security zones, and so on.
- Updating the system software, intrusion rules, or the vulnerability database (VDB).

Keep in mind that you can change some of these configurations from multiple places in the ASA FirePOWER module interface. For example, you can modify security zones using the object manager (**Configuration > ASA FirePOWER Configuration > Object Management**).

Note that the following updates do **not** require policy reapply:

- automatic updates to URL filtering data
- scheduled geolocation database (GeoDB) updates

To determine why an access control or intrusion policy is out of date, use the comparison viewer.

To determine why an access control policy is out of date:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears. Policies that are out of date are marked with red status text that indicates that the ASA FirePOWER module needs a policy update.

Step 2 Click the policy status for an out-of-date policy.

The detailed Apply Access Control Policy pop-up window appears.

Step 3 Click **Out-of-date** next to the changed component you are interested in.

A policy comparison report appears in a new window. For more information, see Comparing Access Control Policies, page 4-16 and Comparing Two Intrusion Policies or Revisions, page 26-9.

Step 4 Optionally, reapply the policy.

See Deploying Configuration Changes.

Deploying Configuration Changes

License: Any

After you use the ASA FirePOWER module to configure your deployment, and any time you make changes to that configuration, you must deploy the new configuration.

This deploy action distributes the following configuration components:

- Access control policies and all associated policies: DNS, file, identity, intrusion, network analysis, SSL
- Any associated rule configurations and objects associated with a policy to be deployed
- Intrusion rule updates
- Device and interface configurations

Caution In special cases, deploying configuration changes may cause a short pause in traffic flow and processing, and may also cause a few packets to pass uninspected. To minimize inconvenience, deploy during a change window.

To deploy configuration changes:

Step 1 Click **Deploy** and select **Deploy FirePOWER Changes**.

Step 2 Click **Deploy**.

Step 3 If the system identifies errors or warnings in the changes to be deployed, you have the following choices:

- Click **Proceed** to continue deploying without resolving error or warning conditions.

- Click **Cancel** to exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.

Troubleshooting Access Control Policies and Rules

License: Any

Properly configuring access control policies, especially creating and ordering access control rules, is a complex task. However, it is a task that is essential to building an effective deployment. If you do not plan your policy carefully, rules can preempt other rules or contain invalid configurations. Both rules and other policy settings can require additional licenses.

To help ensure that the system handles traffic as you expect, the access control policy interface has a robust feedback system. Icons in the access control policy and rule editors mark warnings and errors, as described in the Access Control Error Icons table.

Tip In the access control policy editor, click **Show Warnings** to display a pop-up window that lists all the warnings for the policy.

Additionally, the system warns you at apply-time of any issues that could affect traffic analysis and flow.

Table 4-4 Access Control Error Icons

Icon	Description	Details
	error	If a rule or configuration has an error, you cannot apply the policy until you correct the issue, even if you disable any affected rules.
	warning	You can apply an access control policy that displays rule or other warnings. However, misconfigurations marked with warnings have no effect. For example, you can apply a policy that contains preempted rules or rules that cannot match traffic due to misconfiguration—conditions using empty object groups, application filters that match no applications, configuring URL conditions without having enabled cloud communications, and so on. These rules do not evaluate traffic. If you disable a rule with a warning, the warning icon disappears. It reappears if you enable the rule without correcting the underlying issue. As another example, many features require a specific license. An access control policy successfully applies only to an eligible device.
	information	Information icons convey helpful information about configurations that may affect the flow of traffic. These issues do not prevent you from applying the policy. For example, if you are performing application control or URL filtering, the system may skip matching the first few packets of a connection against some access control rules, until the system identifies the application or web traffic in that connection. This allows connections to be established so that applications and HTTP requests can be identified. For more information, see Limitations to Application Control, page 8-6 and Limitations to URL Detection and Blocking, page 8-11.

Properly configuring access control policies and rules can also reduce the resources required to process network traffic. Creating complex rules, invoking many different intrusion policies, and mis-ordering rules can all affect performance.

For more information, see:

- Simplifying Rules to Improve Performance, page 4-13
- Understanding Rule Preemption and Invalid Configuration Warnings, page 4-14
- Ordering Rules to Improve Performance and Avoid Preemption, page 4-14

Simplifying Rules to Improve Performance

Complex access control policies and rules can command significant resources. When you apply an access control policy, the system evaluates all the rules together and creates an expanded set of criteria that the ASA FirePOWER module uses to evaluate network traffic. A pop-up window may warn that you have exceeded the maximum number of access control rules or intrusion policies supported.

Simplifying Access Control Rules

The following guidelines can help you simplify access control rules and improve performance:

- When constructing a rule, use as few individual elements in your conditions as possible. For example, in network conditions, use IP address blocks rather than individual IP addresses. In port conditions, use port ranges. Use application filters and URL categories and reputations to perform application control and URL filtering, and LDAP user groups to perform user control.

  Note that combining elements into objects that you then use in access control rule conditions does not improve performance. For example, using a network object that contains 50 individual IP addresses gives you only an organizational—not a performance—benefit over including those IP addresses in the condition individually.
- Restrict rules by security zones whenever possible. If a device's interfaces are not in one of the zones in a zone-restricted rule, the rule does not affect performance on that device.
- Do not overconfigure rules. If one condition is enough to match the traffic you want to handle, do not use two.

Avoiding Intrusion Policy and Variable Set Proliferation

The number of unique intrusion policies you can use to inspect traffic in an access control policy depends on the complexity of your policies: you can associate one intrusion policy with each Allow and Interactive Block rule, as well as with the default action. Every unique **pair** of intrusion policy and variable set counts as one policy. You might be able to select as few as three intrusion policies across an entire access control policy.

If you exceed the number of intrusion policies supported, reevaluate your access control policy. You may want to consolidate intrusion policies or variable sets so you can associate a single intrusion policy-variable set pair with multiple access control rules.

Check to see how many policies you select and how many variable sets those policies use in each of the following locations in your access control policy: the **Intrusion Policy used before Access Control rule is determined** option in the Advanced access control policy settings, the default action for the access control policy, and the inspection settings for any access control rules in the policy.

Understanding Rule Preemption and Invalid Configuration Warnings

License: Any

Properly configuring and ordering access control rules (and, in advanced deployments, network analysis rules) is essential to building an effective deployment. Within an access control policy, access control rules can preempt other rules or contain invalid configurations. Similarly, network analysis rules, which you configure using the access control policy's advanced settings, can have the same issues. The system uses warning and error icons to mark these.

Understanding Rule Preemption Warnings

The conditions of an access control rule may preempt a subsequent rule from matching traffic. For example:

```
Rule 1: allow Admin users
Rule 2: block Admin users
```

The second rule above will never block traffic because the first rule will have already allowed the traffic.

Note the following:

- Any type of rule condition can preempt a subsequent rule.
- A rule also preempts an identical subsequent rule where all configured conditions are the same.
- A subsequent rule would not be preempted if any condition is different.

Understanding Invalid Configuration Warnings

Because outside settings that the access control policy depends on may change, an access control policy setting that was valid may become invalid. Consider the following examples:

- If you add a port group to the source ports in a rule, then change the port group to include an ICMP port, the rule becomes invalid and a warning icon appears next to it. You can still apply the policy, but the rule will have no effect on network traffic.
- If you add a user to a rule, then change your LDAP user awareness settings to exclude that user, the rule will have no effect because the user is no longer an access controlled user.

Ordering Rules to Improve Performance and Avoid Preemption

License: Any

Rules in an access control policy are numbered, starting at 1. The system matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Proper access control rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs.

Order Rules from Most to Least Critical

First, you must order rules to suit your organization's needs. Place priority rules that must apply to all traffic near the top of the policy. For example, if you want to inspect traffic from a single user for intrusions (using an Allow rule), but trust all other users in the department (using a Trust rule), place two access control rules in that order.

Order Rules from Specific to General

You can improve performance by placing specific rules earlier, that is, rules that narrowly define the traffic they handle. This is also important because rules with broad conditions can match many different types of traffic, and can preempt later, more specific rules.

Consider a scenario where you want to block most social networking sites, but allow access to certain others. For example, you may want your graphic designers to be able to access Creative Commons Flickr and deviantART content, but not access other sites such as Facebook or Google+. You should order your rules as follows:

```
Rule 1: Allow Flickr, deviantART for the "Design" LDAP user group
Rule 2: Block social networking
```

If you reverse the rules:

```
Rule 1: Block social networking
Rule 2: Allow Flickr, deviantART for the "Design" LDAP user group
```

the first rule blocks all social networking traffic, including Flickr and deviantART. Because no traffic will ever match the second rule, your designers cannot access the content you wanted to make available.

Place Rules that Inspect Traffic Later

Because intrusion, file, and malware inspection require processing resources, placing rules that do not inspect traffic (Trust, Block) before rules that do (Allow, Interactive Block) can improve performance. This is because Trust and Block rules can divert traffic that the system might otherwise have inspected. All other factors being equal, that is, given a set of rules where none is more critical and preemption is not an issue, consider placing them in the following order:

- Monitor rules that log matching connections, but take no other action on traffic
- Trust and Block rules that handle traffic without further inspection
- Allow and Interactive Block rules that do not inspect traffic further
- Allow and Interactive Block rules that optionally inspect traffic for malware, intrusions, or both

Generating a Report of Current Access Control Settings

License: Any

An access control policy report is a record of the policy and rules configuration at a specific point in time. You can use the report, which contains the following information, for auditing purposes or to inspect the current configuration.

Table 4-5 Access Control Policy Report Sections

Section	Description
Policy Information	Provides the name and description of the policy, the name of the user who last modified the policy, and the date and time the policy was last modified.
HTTP Block Response HTTP Interactive Block Response	Provides details on the pages you display to users when you block a website using the policy.
Security Intelligence	Provides details on the policy's Security Intelligence whitelist and blacklist.
Default Action	Lists the default action and associated variable set, if any.
Rules	Lists each access control rule in the policy, and provides details about its configuration.

Table 4-5 Access Control Policy Report Sections (continued)

Section	Description
Advanced Settings	Detailed information on the policy's advanced settings, including: • network analysis policies used to preprocess traffic for the access control policy, as well as global preprocessing options • adaptive profile settings for passive deployments • performance settings for detecting files, malware, and intrusions • other policy-wide settings
Referenced Objects	Provides details on the reusable objects referenced by the access control policy, including intrusion policy variable sets and objects used by the SSL policy.

You can also generate an access control comparison report that compares a policy with the currently applied policy or with another policy. For more information, see Comparing Access Control Policies, page 4-16.

To view an access control policy report:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the report icon ( ) next to the policy for which you want to generate a report. Remember to save any changes before you generate an access control policy report; only saved changes appear in the report.

The system generates the report. You are prompted to save the report to your computer.

Comparing Access Control Policies

License: Any

To review policy changes for compliance with your organization's standards or to optimize system performance, you can examine the differences between two access control policies. You can compare any two policies or the currently applied policy with another policy. Optionally, after you compare, you can then generate a PDF report to record the differences between the two policies.

There are two tools you can use to compare policies:

- The comparison view displays only the differences between two policies in a side-by-side format. The name of each policy appears in the title bar on the left and right sides of the comparison view except when you select **Running Configuration**, in which case a blank bar represents the currently active policy.

  You can use this to view and navigate both policies on the module interface, with their differences highlighted.

- The comparison report creates a record of only the differences between two policies in a format similar to the policy report, but in PDF format.

  You can use this to save, copy, print, and share your policy comparisons for further examination.

For more information on understanding and using the policy comparison tools, see:

- Using the Access Control Policy Comparison View, page 4-17
- Using the Access Control Policy Comparison Report, page 4-17

Using the Access Control Policy Comparison View

License: Any

The comparison view displays both policies in a side-by-side format, with each policy identified by name in the title bar on the left and right sides of the comparison view. When comparing two policies other than the running configuration, the time of last modification and the last user to modify are displayed with the policy name.

Differences between the two policies are highlighted:

- Blue indicates that the highlighted setting is different in the two policies, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one policy but not the other.

You can perform any of the actions in the following table.

Table 4-6 Access Control Policy Comparison View Actions

To...	You can...
navigate individually through changes	click **Previous** or **Next** above the title bar. The double-arrow icon ( ) centered between the left and right sides moves, and the **Difference** number adjusts to identify which difference you are viewing.
generate a new policy comparison view	click **New Comparison.** The **Select Comparison** window appears. See Using the Access Control Policy Comparison Report, page 4-17 for more information.
generate a policy comparison report	click **Comparison Report**. The policy comparison report creates a PDF document that lists only the differences between the two policies.

Using the Access Control Policy Comparison Report

License: Any

An access control policy comparison report is a record of all differences between two access control policies or a policy and the currently applied policy identified by the policy comparison view, presented in PDF format. You can use this report to further examine the differences between two policy configurations and to save and disseminate your findings.

You can generate an access control policy comparison report from the comparison view for any policies to which you have access. Remember to save any changes before you generate a policy report; only saved changes appear in the report.

The format of the policy comparison report is the same as the policy report with one exception: the policy report contains all configurations in the policy, and the policy comparison report lists only those configurations that differ between the policies. An access control policy comparison report contains the sections described in Table 4-5 on page 4-15.

Tip You can use a similar procedure to compare SSL, network analysis, intrusion, file, or system policies.

To compare two access control policies:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click **Compare Policies**.

The Select Comparison window appears.

Step 3 From the **Compare Against** drop-down list, select the type of comparison you want to make:

- To compare two different policies, select **Other Policy**.

  The page refreshes and the Policy A and Policy B drop-down lists appear.
- To compare another policy to the currently active policy, select **Running Configuration.**

  The page refreshes and the Target/Running Configuration A and Policy B drop-down lists appear.

Step 4 Depending on the comparison type you selected, you have the following choices:

- If you are comparing two different policies, select the policies you want to compare from the Policy A and Policy B drop-down lists.
- If you are comparing the running configuration to another policy, select the second policy from the Policy B drop-down list.

Step 5 Click **OK** to display the policy comparison view.

The comparison view appears.

Step 6 Optionally, click **Comparison Report** to generate the access control policy comparison report.

The access control policy comparison report appears. You are prompted to save the report to your computer.



Blacklisting Using Security Intelligence IP Address Reputation

As a first line of defense against malicious Internet content, the ASA FirePOWER module includes the Security Intelligence feature, which allows you to immediately blacklist (block) connections based on the latest reputation intelligence, removing the need for a more resource-intensive, in-depth analysis. Security Intelligence filtering requires a Protection license.

Security Intelligence works by blocking traffic to or from IP addresses that have a known bad reputation. This traffic filtering takes place **before** any other policy-based inspection, analysis, or traffic handling.

Note that you could create access control rules that perform a similar function to Security Intelligence filtering by manually restricting traffic by IP address. However, access control rules are wider in scope, more complex to configure, and cannot automatically update using dynamic feeds.

Traffic blacklisted by Security Intelligence is immediately blocked and therefore is not subject to any further inspection—not for intrusions, exploits, malware, and so on. Optionally, and recommended in passive deployments, you can use a "monitor-only" setting for Security Intelligence filtering. This allows the system to analyze connections that would have been blacklisted, but also logs the match to the blacklist and generates an end-of-connection security intelligence event.

For your convenience, Cisco provides the *Intelligence Feed* (sometimes called the *Sourcefire Intelligence Feed*), which is comprised of several regularly updated collections of IP addresses determined by the VRT to have a poor reputation. The Intelligence Feed tracks open relays, known attackers, bogus IP addresses (bogon), and so on. You can also customize the feature to suit the unique needs of your organization, for example:

- **third-party feeds**—you can supplement the Intelligence Feed with third-party reputation feeds, which the system can automatically update just as it does the Cisco feed
- **custom blacklists**—the system allows you to manually blacklist specific IP addresses in many ways depending on your needs
- **enforcing blacklisting by security zone**—to improve performance, you may want to target enforcement, for example, restricting spam blacklisting to a zone that handles email traffic
- **monitoring instead of blacklisting**—especially useful in passive deployments and for testing feeds before you implement them; you can merely monitor the violating sessions instead of blocking them, generating end-of-connection events
- **whitelisting to eliminate false positives**—when a blacklist is too broad in scope, or incorrectly blocks traffic that you want to allow (for example, to vital resources), you can override a blacklist with a custom whitelist

For detailed information on configuring your access control policy to perform Security Intelligence to perform Security Intelligence filtering and viewing the event data that this filtering produces, see the following sections:

- Choosing a Security Intelligence Strategy, page 5-2
- Building the Security Intelligence Whitelist and Blacklist, page 5-3
- Logging Security Intelligence (Blacklisting) Decisions, page 36-8

Choosing a Security Intelligence Strategy

License: Protection

The easiest way to construct a blacklist is to use the Intelligence Feed, which tracks IP addresses known to be open relays, known attackers, bogus IP addresses (bogon), and so on. Because the Intelligence Feed is regularly updated, using it ensures that the system uses up-to-date information to filter your network traffic. Malicious IP addresses that represent security threats such as malware, spam, botnets, and phishing may appear and disappear faster than you can update and apply new policies.

To augment the Intelligence Feed, you can perform Security Intelligence filtering using custom or third-party IP address lists and feeds, where:

- a *list* is a static list of IP addresses that you upload to the ASA FirePOWER module
- a *feed* is a dynamic list of IP addresses that the ASA FirePOWER module downloads from the Internet on a regular basis; the Intelligence Feed is a special kind of feed

For detailed information on configuring Security Intelligence lists and feeds, including Internet access requirements, see Working with Security Intelligence Lists and Feeds, page 2-4.

Using the Security Intelligence Global Blacklist

In the course of your analysis, you can build a *global blacklist*. For example, if you notice a set of routable IP addresses in intrusion events associated with exploit attempts, you can blacklist those IP addresses. The ASA FirePOWER module uses this global blacklist (and a related *global whitelist*) to perform Security Intelligence filtering in all access control policies. For information on managing these global lists, see Working with the Global Whitelist and Blacklist, page 2-6.

Note Although feed updates and additions to the global blacklist (or global whitelist; see below) automatically implement changes throughout your deployment, any other change to a Security Intelligence object requires you to reapply the access control policy. For more information, see Table 2-1 on page 2-5.

Using Network Objects

Finally, a simple way to construct a blacklist is to use *network objects or network object groups* that represent an IP address, IP address block, or collection of IP addresses. For information on creating and modifying network objects, see Working with Network Objects, page 2-3.

Using Security Intelligence Whitelists

In addition to a blacklist, each access control policy has an associated whitelist, which you can also populate with Security Intelligence objects. A policy's whitelist overrides its blacklist. That is, the system evaluates traffic with a whitelisted source or destination IP address using access control rules, even if the IP address is also blacklisted. In general, use the whitelist if a blacklist is still useful, but is too broad in scope and incorrectly blocks traffic that you want to inspect.

For example, if a reputable feed improperly blocks your access to vital resources but is overall useful to your organization, you can whitelist only the improperly classified IP addresses, rather than removing the whole feed from the blacklist.

Enforcing Security Intelligence Filtering by Security Zone

For added granularity, you can enforce Security Intelligence filtering based on whether the source or destination IP address in a connection resides in a particular security zone.

To extend the whitelist example above, you could whitelist the improperly classified IP addresses, but then restrict the whitelist object using a security zone used by those in your organization who need to access those IP addresses. That way, only those with a business need can access the whitelisted IP addresses. As another example, you could use a third-party spam feed to blacklist traffic on an email server security zone.

Monitoring—Rather than Blacklisting—Connections

If you are not sure whether you want to blacklist a particular IP address or set of addresses, you can use a "monitor-only" setting, which allows the system to pass the matching connection to access control rules, but also logs the match to the blacklist and generates an end-of-connection Security Intelligence event. Note that you cannot set the global blacklist to monitor-only.

Consider a scenario where you want to test a third-party feed before you implement blocking using that feed. When you set the feed to monitor-only, the system allows connections that would have been blocked to be further analyzed by the system, but also logs a record of each of those connections for your evaluation.

In passive deployments, to optimize performance, Cisco recommends that you always use monitor-only settings. Devices that are deployed passively cannot affect traffic flow; there is no advantage to configuring the system to block traffic. Additionally, because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.

Building the Security Intelligence Whitelist and Blacklist

License: Protection

To build a whitelist and blacklist, you populate them with any combination of network objects and groups, as well as Security Intelligence feeds and lists, all of which you can constrain by security zone.

By default, access control policies use the ASA FirePOWER module's global whitelist and blacklist, which apply to any zone. These lists are populated by your analysts. You can opt not to use these global lists on a per-policy basis.



Note You cannot apply an access control policy that uses a populated global whitelist or blacklist to a device not licensed for Protection. If you added IP addresses to either global list, you **must** remove the non-empty list from the policy's Security Intelligence configuration before you can apply the policy. For more information, see Working with the Global Whitelist and Blacklist, page 2-6.

After you build your whitelist and blacklist, you can log blacklisted connections. You can also set individual blacklisted objects, including feeds and lists, to monitor-only. This allows the system to handle connections involving blacklisted IP addresses using access control, but also logs the connection's match to the blacklist.

Use the Security Intelligence tab in the access control policy to configure the whitelist, blacklist, and logging options. The page lists the Available Objects you can use in either the whitelist or blacklist, as well as the Available Zones you can use to constrain whitelisted and blacklisted objects. Each type of object or zone is distinguished with an different icon. The objects marked with the Cisco icon ( ) represent the different categories in the Intelligence Feed.

In the blacklist, objects set to block are marked with the block icon ( ) while monitor-only objects are marked with the monitor icon ( ). Because the whitelist overrides the blacklist, if you add the same object to both lists, the system displays the blacklisted object with a strikethrough.

You can add up to a total of 255 objects to the whitelist and the blacklist. That is, the number of objects in the whitelist plus the number in the blacklist cannot exceed 255.

Note that although you can add network objects with a netmask of `/0` to the whitelist or blacklist, address blocks using a `/0` netmask in those objects will be ignored and whitelist and blacklist filtering will not occur based on those addresses. Address blocks with a `/0` netmask from Security Intelligence feeds are also ignored. If you want to monitor or block all traffic targeted by a policy, use an access control rule with the **Monitor** or **Block** rule action, respectively, and a default value of **any** for the **Source Networks** and **Destination Networks**, instead of Security Intelligence filtering.

To build the Security Intelligence whitelist and blacklist for an access control policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to configure.

The access control policy editor appears.

Step 3 Select the **Security Intelligence** tab.

Security Intelligence settings for the access control policy appear.

Step 4 Optionally, click the logging icon ( ) to log blacklisted connections.

You must enable logging before you can set blacklisted objects to monitor-only. For details, see Logging Security Intelligence (Blacklisting) Decisions, page 36-8.

Step 5 Begin building your whitelist and blacklist by selecting one or more **Available Objects**.

Use Shift and Ctrl to select multiple objects, or right-click and **Select All**.

Tip You can search for existing objects to include, or create objects on the fly if no existing objects meet the needs of your organization. For more information, see Searching for Objects to Whitelist or Blacklist, page 5-5.

Step 6 Optionally, constrain the selected objects by zone by selecting an **Available Zone**.

By default, objects are not constrained, that is, they have a zone of `Any`. Note that other than using **Any**, you can constrain by only one zone. To enforce Security Intelligence filtering for an object on multiple zones, you must add the object to the whitelist or blacklist separately for each zone. Also, the global whitelist or blacklist cannot be constrained by zone.

Step 7 Click **Add to Whitelist** or **Add to Blacklist**.

You can also click and drag the selected objects to either list.

The objects you selected are added to the whitelist or blacklist.

Tip To remove an object from a list, click its delete icon ( ). Use Shift and Ctrl to select multiple objects, or right-click and **Select All**, then right-click and select **Delete Selected**. If you are deleting a global list, you must confirm your choice. Note that removing an object from a whitelist or blacklist does not delete that object from the ASA FirePOWER module.

Step 8 Repeat steps 5 through 7 until you are finished adding objects to your whitelist and blacklist.

Step 9 Optionally, set blacklisted objects to monitor-only by right-clicking the object under **Blacklist**, then selecting **Monitor-only (do not block)**.

In passive deployments, Cisco recommends you set all blacklisted objects to monitor-only. Note, however, that you cannot set the global blacklist to monitor-only.

Step 10 Click **Store ASA FirePOWER Changes**.

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Searching for Objects to Whitelist or Blacklist

License: Protection

If you have multiple network objects, groups, feeds, and lists, use the search feature to narrow the objects you want to blacklist or whitelist.

To search for objects to whitelist or blacklist:

Step 1 Type your query in the **Search by name or value** field.

The Available Objects list updates as you type to display matching items. To clear the search string, click the reload icon ( ) above the search field or click the clear icon ( ) in the search field.

You can search on network object names and on the values configured for those objects. For example, if you have an individual network object named `Texas Office` with the configured value `192.168.3.0/24`, and the object is included in the group object `US Offices`, you can display both objects by typing a partial or complete search string such as `Tex`, or by typing a value such as `3`.



Tuning Traffic Flow Using Access Control Rules

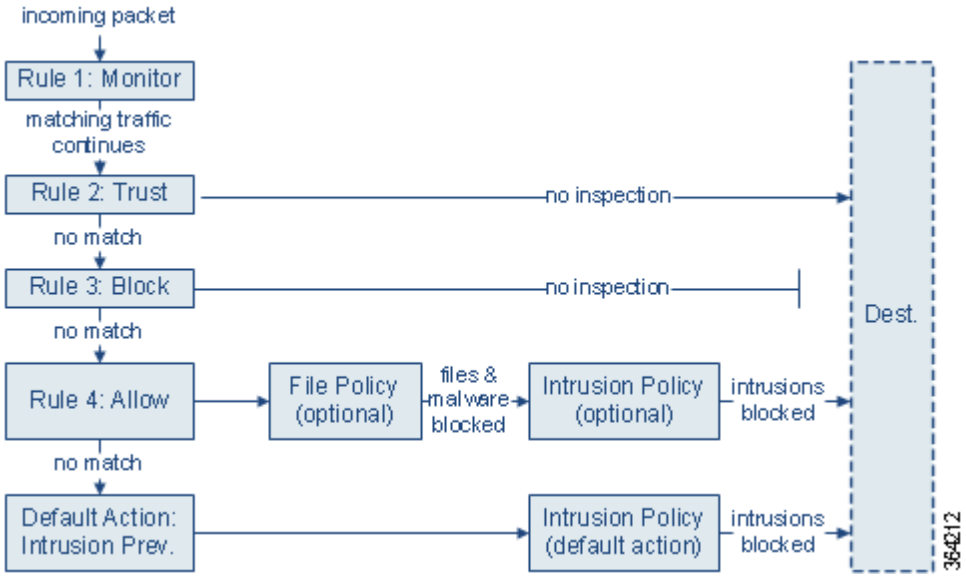
Within an access control policy, *access control rules* provide a granular method of handling network traffic.

Note Security Intelligence-based traffic filtering, and some decoding and preprocessing occur **before** network traffic is evaluated by access control rules. You can also configure the *SSL inspection* feature to block or decrypt encrypted traffic before access control rules evaluate it.

The system matches traffic to access control rules in the order you specify. In most cases, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic. Conditions can be simple or complex; you can control traffic by security zone, network or geographical location, port, application, requested URL, and user.

Each rule also has an *action*, which determines whether you monitor, trust, block, or allow matching traffic. When you allow traffic, you can specify that the system first inspect it with intrusion or file policies to block any exploits, malware, or prohibited files before they reach your assets or exit your network. However, after the system trusts or blocks traffic, it does **not** perform further inspection.

The following scenario summarizes the ways that traffic can be evaluated by access control rules in an inline, intrusion prevention deployment.



In this scenario, traffic is evaluated as follows:

- **Rule 1: Monitor** evaluates traffic first. Monitor rules track and log network traffic but do not affect traffic flow. The system continues to match traffic against additional rules to determine whether to permit or deny it.
- **Rule 2: Trust** evaluates traffic next. Matching traffic is allowed to pass to its destination without further inspection. Traffic that does not match continues to the next rule.
- **Rule 3: Block** evaluates traffic third. Matching traffic is blocked without further inspection. Traffic that does not match continues to the final rule.
- **Rule 4: Allow** is the final rule. For this rule, matching traffic is allowed; however, prohibited files, malware, intrusions, and exploits within that traffic are detected and blocked. Remaining non-prohibited, non-malicious traffic is allowed to its destination. Note that you might have additional Allow rules that perform only file inspection, or only intrusion inspection, or neither.
- **Default Action** handles all traffic that does not match any of the rules. In this scenario, the default action performs intrusion prevention before allowing non-malicious traffic to pass. In a different deployment, you might have a default action that trusts or blocks all traffic, without further inspection. (You cannot perform file or malware inspection on traffic handled by the default action.)

For more information on access control rules, see:

- Creating and Editing Access Control Rules, page 6-2
- Managing Access Control Rules in a Policy, page 6-11
- Troubleshooting Access Control Policies and Rules, page 4-12

Creating and Editing Access Control Rules

License: Any

Within an access control policy, access control rules provide a granular method of handling network traffic. In addition to its unique name, each access control rule has the following basic components:

State

By default, rules are enabled. If you disable a rule, the system does not use it to evaluate network traffic, and stops generating warnings and errors for that rule.

Position

Rules in an access control policy are numbered, starting at 1. The system matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Conditions

Conditions specify the specific traffic the rule handles. Conditions can match traffic by security zone, network or geographical location, port, application, requested URL, or user. Conditions can be simple or complex; their use often depends on license.

Action

A rule's action determines how the system handles matching traffic. You can monitor, trust, block, or allow (with or without further inspection) matching traffic. Note that the system does **not** perform inspection on trusted or blocked traffic.

Inspection

Inspection options for an access control rule govern how the system inspects and blocks malicious traffic you would otherwise allow. When you allow traffic with a rule, you can specify that the system first inspect it with intrusion or file policies to block any exploits, malware, or prohibited files before they reach your assets or exit your network.

Logging

A rule's logging settings govern the records the system keeps of the traffic it handles. You can keep a record of traffic that matches a rule. In general, you can log sessions at the beginning and end of a connection. You can log connections to the ASA FirePOWER module, as well as to the system log (syslog) or to an SNMP trap server.

Comments

Each time you save changes to an access control rule, you can add a comment.

Use the access control rule editor to add and edit access control rules; access the rule editor from the Rules tab of the access control policy editor. In the rule editor, you:

- Configure basic properties such as the rule's name, state, position, and action in the upper portion of the editor.
- Add conditions using the tabs on the left side of the lower portion of the editor.
- Use the tabs on the right side of the lower portion to configure inspection and logging options, and also to add comments to the rule. For your convenience, the editor lists the rule's inspection and logging options regardless of which tab you are viewing.



Note Properly creating and ordering access control rules is a complex task, but one that is essential to building an effective deployment. If you do not plan your policy carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. To help ensure that the system handles traffic as

you expect, the access control policy interface has a robust warning and error feedback system for rules. For more information, see Troubleshooting Access Control Policies and Rules, page 4-12.

To create or modify an access control rule:

Step 1 Choose **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

Step 2 Click the edit icon ( ) next to the access control policy where you want to add a rule.

Step 3 You have the following options:

- To add a new rule, click **Add Rule**.
- To edit an existing rule, click the edit icon ( ) next to the rule you want to edit.

Step 4 Enter a **Name** for the rule.

Each rule must have a unique name. You can use up to thirty printable characters, including spaces and special characters, with the exception of the colon (:).

Step 5 Configure the rule components, as summarized above. You can configure the following, or accept the defaults:

- Specify whether the rule is **Enabled**.
- Specify the rule position; see Specifying a Rule's Order of Evaluation, page 6-4.
- Specify a rule **Action**; see Using Rule Actions to Determine Traffic Handling and Inspection, page 6-7.
- Configure the rule's conditions; see Using Conditions to Specify the Traffic a Rule Handles, page 6-5.
- For Allow and Interactive Block rules, configure the rule's **Inspection** options; see Controlling Traffic Using Intrusion and File Policies, page 11-1.
- Configure content restriction settings by clicking the Safe Search ( ) or YouTube EDU ( ) icon on the **Applications** tab. If the icons are dimmed, content restriction is disabled for the rule. For more information, see Using Access Control Rules to Enforce Content Restriction, page 13-1.
- Specify **Logging** options; see Logging Connections in Network Traffic, page 36-1.
- Add **Comments**; see Adding Comments to a Rule, page 6-10.

Step 6 Click **Store FirePOWER Changes** to save the rule.

Your rule is saved. You can click the delete icon ( ) to delete the rule. You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Specifying a Rule's Order of Evaluation

License: Any

When you first create an access control rule, you specify its position using the **Insert** drop-down list in the rule editor. Rules in an access control policy are numbered, starting at 1. The system matches traffic to access control rules in top-down order by ascending rule number.

In most cases, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic. Except in the case of Monitor rules (which log traffic but do not affect traffic flow), the system does **not** continue to evaluate traffic against additional, lower-priority rules after that traffic matches a rule.

Tip Proper access control rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs. For more information, see Ordering Rules to Improve Performance and Avoid Preemption, page 4-14.

In addition to ordering rules by number, you can group rules by category. By default the system provides three categories: Administrator, Standard, and Root. You can add custom categories, but you cannot delete the Cisco-provided categories or change their order. For information on changing the position or category of an existing rule, see Changing a Rule's Position or Category, page 6-13.

To add a rule to a category while editing or creating a rule:

Step 1 In the access control rule editor, from the **Insert** drop-down list, select **Into Category**, then select the category you want to use.

When you save the rule, it is placed last in that category.

To position a rule by number while editing or creating a rule:

Step 1 In the access control rule editor, from the **Insert** drop-down list, select **above rule** or **below rule**, then type the appropriate rule number.

When you save the rule, it is placed where you specified.

Using Conditions to Specify the Traffic a Rule Handles

License: feature dependent

An access control rule's conditions identify the type of traffic that rule handles. Conditions can be simple or complex; you can control traffic by security zone, network or geographical location, port, application, requested URL, and user.

When adding conditions to access control rules, keep the following points in mind:

- You can configure multiple conditions per rule. Traffic must match **all** the conditions in the rule for the rule to apply to traffic. For example, you can use a single rule to perform URL filtering (URL condition) for specific hosts (zone or network condition).
- For each condition in a rule, you can add up to 50 criteria. Traffic that matches **any** of a condition's criteria satisfies the condition. For example, you can use a single rule to perform user control for up to 50 users and groups.

Note that you can constrain zone and network conditions by source and destination, using up to 50 source and up to 50 destination criteria. If you add both source and destination criteria to a zone or network condition, matching traffic must originate from one of the specified source zones/networks **and** egress through one of the destination zones/networks. In other words, the system links multiple condition criteria of the same type with an OR operation, and links multiple condition types with an AND operation. For example, if your rule conditions are:

```
Source Networks: 10.0.0.0/8, 192.168.0.0/16
Application Category: peer to peer
```

the rule would match peer-to-peer application traffic from a host on one of your private IPv4 networks—a packet must originate from either one **OR** the other source network, **AND** represent peer-to-peer application traffic. Both of the following connections trigger the rule:

```
10.42.0.105 to anywhere, using LimeWire
192.168.42.105 to anywhere, using Kazaa
```

If you do not configure a particular condition for a rule, the system does not match traffic based on that criterion. For example, a rule with a network condition but no application condition evaluates traffic based on its source or destination, regardless of the application used in the session.

Note When you apply an access control policy, the system evaluates all its rules and creates an expanded set of criteria that the ASA FirePOWER module uses to evaluate network traffic. Complex access control policies and rules can command significant resources. For tips on simplifying access control rules and other ways to improve performance, see Troubleshooting Access Control Policies and Rules, page 4-12.

When you add or edit an access control rule, use the tabs on the left side of the lower portion of the rule editor to add and edit rule conditions. The following table summarizes the types of conditions you can add.

Table 6-1 Access Control Rule Condition Types

These Conditions...	Match Traffic...	Details
Zones	entering or leaving a device via an interface in a specific security zone	A security zone is a logical grouping of one or more interfaces according to your deployment and security policies. To build a zone condition, see Controlling Traffic by Security Zone, page 7-1.
Networks	by its source or destination IP address, country, or continent	You can explicitly specify IP addresses or address blocks. The geolocation feature also allows you to control traffic based on its source or destination country or continent. To build a network condition, see Controlling Traffic by Network or Geographical Location, page 7-3.
Ports	by its source or destination port	For TCP and UDP, you can control traffic based on the transport layer protocol. For ICMP and ICMPv6 (IPv6-ICMP), you can control traffic based on its Internet layer protocol plus an optional type and code. Using port conditions, you can also control traffic using other protocols that do not use ports. To build a port condition, see Controlling Traffic by Port and ICMP Codes, page 7-5.
Applications	by the application detected in a session	You can control access to individual applications, or filter access according to basic characteristics: type, risk, business relevance, categories, and tags. To build an application condition, see Controlling Application Traffic, page 8-2.

Table 6-1 Access Control Rule Condition Types (continued)

These Conditions...	Match Traffic...	Details
URLs	by the URL requested in the session	You can limit the websites that users on your network can access either individually or based on the URL's general classification and risk level. To build a URL condition, see Blocking URLs, page 8-7.
Users	by the user involved in the session	You can control traffic based on the LDAP user logged into a host involved in a monitored session. You can control traffic based on individual users or groups retrieved from a Microsoft Active Directory server. To build a user condition, see Access Control Rules: Realms and Users, page 9-1.

Note that although you can create access control rules with any license, certain rule conditions require that you enable specific licensed capabilities before you can apply the policy. For more information, see License Requirements for Access Control, page 4-2.

Using Rule Actions to Determine Traffic Handling and Inspection

License: Any

Every access control rule has an *action* that determines the following for matching traffic:

- handling—foremost, the rule action governs whether the system will monitor, trust, block, or allow traffic that matches the rule's conditions
- inspection—certain rule actions allow you, when properly licensed, to further inspect matching traffic before allowing it to pass
- logging—the rule action determines when and how you can log details about matching traffic

The access control policy's *default action* handles traffic that does not meet the conditions of any non-Monitor access control rule; see Setting Default Handling and Inspection for Network Traffic, page 4-4.

Keep in mind that only devices deployed inline can block or modify traffic. Devices deployed passively can analyze and log, but not affect, the flow of traffic. For detailed information on rule actions and how they affect traffic handling, inspection, and logging, see the following sections:

- Monitor Action: Postponing Action and Ensuring Logging, page 6-8
- Trust Action: Passing Traffic Without Inspection, page 6-8
- Blocking Actions: Blocking Traffic Without Inspection, page 6-8
- Interactive Blocking Actions: Allowing Users to Bypass Website Blocks, page 6-9
- Allow Action: Allowing and Inspecting Traffic, page 6-10
- Controlling Traffic Using Intrusion and File Policies, page 11-1
- Logging Connections Based on Access Control Handling, page 36-9

Monitor Action: Postponing Action and Ensuring Logging

License: Any

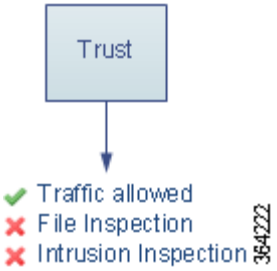
The **Monitor** action does not affect traffic flow; matching traffic is neither immediately permitted nor denied. Rather, traffic is matched against additional rules to determine whether to permit or deny it. The first non-Monitor rule matched determines traffic flow and any further inspection. If there are no additional matching rules, the system uses the default action.

Because the primary purpose of Monitor rules is to track network traffic, the system automatically logs end-of connection events for monitored traffic. That is, connections are logged even if the traffic matches no other rules and you do not enable logging on the default action. For more information, see Understanding Logging for Monitored Connections, page 36-5.

Trust Action: Passing Traffic Without Inspection

License: Any

The **Trust** action allows traffic to pass without further inspection of any kind.

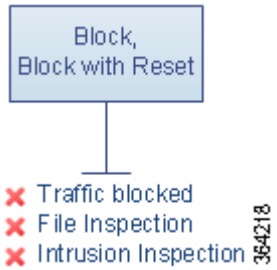


You can log trusted network traffic at both the beginning and end of connections. For more information, see Understanding Logging for Trusted Connections, page 36-5.

Blocking Actions: Blocking Traffic Without Inspection

License: Any

The **Block** and **Block with reset** actions deny traffic without further inspection of any kind. Block with reset rules also reset the connection.



For decrypted HTTP traffic, when the system blocks a web request, you can override the default browser or server page with a custom page that explains that the connection was denied. The system calls this custom page an *HTTP response page*; see Displaying a Custom Web Page for Blocked URLs, page 8-14.

You can log blocked network traffic only at the beginning of connections. Note that only devices deployed inline can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection. For more information, see Understanding Logging for Blocked and Interactively Blocked Connections, page 36-5.

Caution Logging blocked TCP connections during a Denial of Service (DoS) attack can affect system performance with multiple similar events. Before you enable logging for an Block rule, consider whether the rule monitors traffic on an Internet-facing interface or other interface vulnerable to DoS attack.

Interactive Blocking Actions: Allowing Users to Bypass Website Blocks

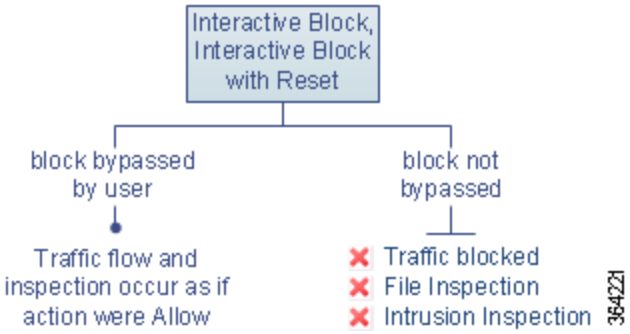
License: Any

For decrypted HTTP traffic, the **Interactive Block** and **Interactive Block with reset** actions give users a chance to bypass a website block by clicking through a customizable warning page, called an *HTTP response page*. Interactive Block with reset rules also reset the connection.

If you configure SSL inspection to decrypt web traffic and that traffic matches an Interactive Block rule, the system encrypts the response page and sends it at the end of the reencrypted SSL response stream.

For all interactively blocked traffic, the system's handling, inspection, and logging depend on whether the user bypasses the block:

- If a user does not (or cannot) bypass the block, the rule mimics a Block rule. Matching traffic is denied without further inspection and you can log only the beginning of the connection. These beginning-of-connection events have an `Interactive Block` or `Interactive Block with Reset` action.
- If a user bypasses the block, the rule mimics an Allow rule. Therefore, you can associate either type of Interactive Block rule with a file and intrusion policy to inspect this user-allowed traffic. The system can also log both beginning and end-of-connection events. These connection events have an action of `Allow`.



Allow Action: Allowing and Inspecting Traffic

License: Any

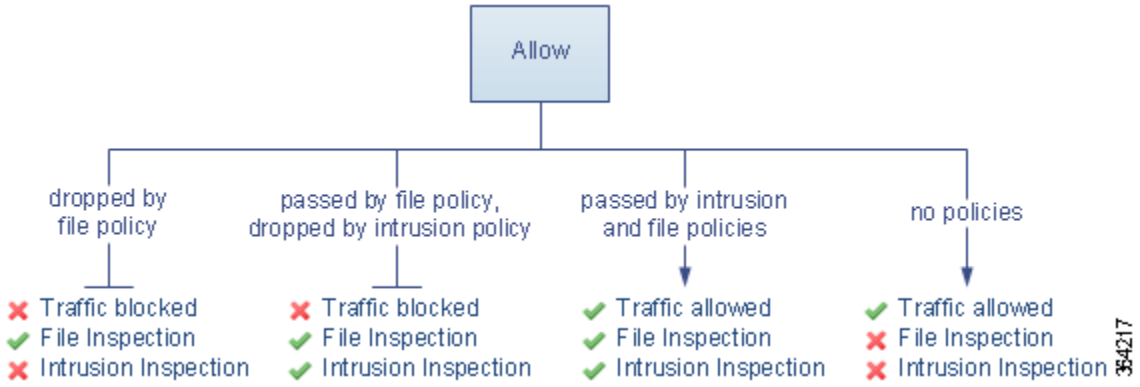
The **Allow** action allows matching traffic to pass. When you allow traffic, you can use an associated intrusion or file policy (or both) to further inspect and block unencrypted or decrypted network traffic:

- With a Protection license, you can use an intrusion policy to analyze network traffic according to intrusion detection and prevention configurations and, optionally, drop offending packets.
- Also with a Protection license, you can perform file control using a file policy. File control allows you to detect and block your users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols.
- With a Malware license, you can perform network-based advanced malware protection (AMP), also using a file policy. Network-based AMP can inspect files for malware, and optionally block detected malware.

For instructions on how to associate an intrusion or file policy with an access control rule, see Controlling Traffic Using Intrusion and File Policies, page 11-1.

The diagram below illustrates the types of inspection performed on traffic that meets the conditions of an Allow rule (or a user-bypassed Interactive Block rule; see Interactive Blocking Actions: Allowing Users to Bypass Website Blocks, page 6-9). Notice that file inspection occurs before intrusion inspection; blocked files are not inspected for intrusion-related exploits.

For simplicity, the diagram displays traffic flow for situations where both (or neither) an intrusion and a file policy are associated with an access control rule. You can, however, configure one without the other. Without a file policy, traffic flow is determined by the intrusion policy; without an intrusion policy, traffic flow is determined by the file policy.



You can log allowed network traffic at both the beginning and end of connections.

Adding Comments to a Rule

License: Any

When you create or edit an access control rule, you can add a comment. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change. You can display a list of all comments for a rule along with the user who added each comment and the date the comment was added.

When you save a rule, all comments made since the last save become read-only.

To add a comment to a rule:

Step 1 In the access control rule editor, select the **Comments** tab.

The Comments page appears.

Step 2 Click **New Comment**.

The New Comment pop-up window appears.

Step 3 Type your comment and click **OK**.

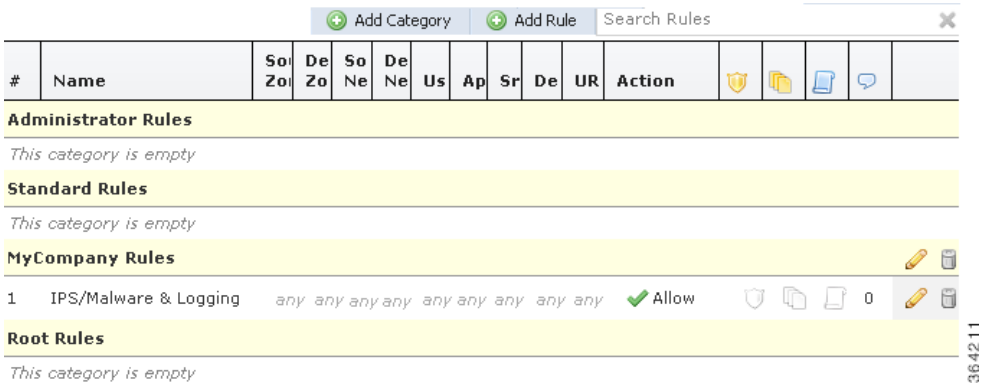
Your comment is saved. You can edit or delete this comment until you save the rule.

Step 4 Save or continue editing the rule.

Managing Access Control Rules in a Policy

License: Any

The Rules tab of the access control policy editor, shown in the following graphic, allows you to add, edit, search, move, enable, disable, delete, and otherwise manage access control rules within your policy.



For each rule, the policy editor displays its name, a summary of its conditions, the rule action, plus icons that communicate the rule's inspection and logging options. Other icons represent comments, warnings, errors, and other important information, as described in the following table. Disabled rules are grayed and marked `(disabled)` beneath the rule name.

Table 6-2 Understanding the Access Control Policy Editor

Icon	Description	You can...
	intrusion inspection	Click an active (yellow) inspection icon to edit the inspection options for the rule; see Controlling Traffic Using Intrusion and File Policies, page 11-1. If the icon is inactive (white), no policy of that type is selected for the rule.
	file and malware inspection	

Table 6-2 Understanding the Access Control Policy Editor (continued)

Icon	Description	You can...
	logging	Click an active (blue) logging icon to edit the logging options for the rule; see Logging Connections Based on Access Control Handling, page 36-9. If the icon is inactive (white), connection logging is disabled for the rule.
	comment	Click the number in the comment column to add a comment to a rule; see Adding Comments to a Rule, page 6-10. The number indicates how many comments the rule already contains.
	warning	In the access control policy editor, click **Show Warnings** to display a pop-up window that lists all the warnings for the policy; see Troubleshooting Access Control Policies and Rules, page 4-12.
	error	
	information	

For information on managing access control rules, see:

- Creating and Editing Access Control Rules, page 6-2
- Searching Access Control Rules, page 6-12
- Enabling and Disabling Rules, page 6-13
- Changing a Rule's Position or Category, page 6-13

Searching Access Control Rules

License: Any

You can search the list of access control rules for matching values using an alphanumeric string, including spaces and printable, special characters. The search inspects the rule name and any rule condition you have added to the rule. For rule conditions, the search matches any name or value you can add for each condition type (zone, network, application, and so on). This includes individual object names or values, group object names, individual object names or values within a group, and literal values.

You can use complete or partial search strings. The column for matching values is highlighted for each matching rule. For example, if you search on all or part of the string `100Bao`, at a minimum, the Applications column is highlighted for each rule where you have added the 100Bao application. If you also have a rule named 100Bao, both the Name and Applications columns are highlighted.

You can navigate to each previous or next matching rule. A status message displays the current match and the total number of matches.

Matches may occur on any page of a multi-page rule list. When the first match is not on the first page, the page where the first match occurs is displayed. Selecting the next match when you are at the last match takes you to the first match, and selecting the previous match when you are at the first match takes you to the last match.

To search for rules:

Step 1 In the access control policy editor for the policy you want to search, click the **Search Rules** prompt, type a search string, then press Enter. You can also use the Tab key or click a blank page area to initiate the search.

Columns for rules with matching values are highlighted, with differentiated highlighting for the indicated (first) match.

Step 2 Find the rules you are interested in:

- To navigate between matching rules, click the next-match (▼) or previous-match (▲) icon.
- To refresh the page and clear the search string and any highlighting, click the clear icon (✖).

Enabling and Disabling Rules

License: Any

When you create an access control rule, it is enabled by default. If you disable a rule, the system does not use it to evaluate network traffic and stops generating warnings and errors for that rule. When viewing the list of rules in an access control policy, disabled rules are grayed out, although you can still modify them. Note that you can also enable or disable an access control rule using the rule editor; see Creating and Editing Access Control Rules, page 6-2.

To change an access control rule's state:

Step 1 In the access control policy editor for the policy that contains the rule you want to enable or disable, right-click the rule and choose a rule state:

- To enable an inactive rule, select **State > Enable**.
- To disable an active rule, select **State > Disable**.

Step 2 Click **Store FirePOWER Changes** to save the policy.

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Changing a Rule's Position or Category

License: Any

To help you organize access control rules, every access control policy has three system-provided rule categories: Administrator Rules, Standard Rules, and Root Rules. You cannot move, delete, or rename these categories, although you can create custom categories.

For more information, see:

- Moving a Rule, page 6-14
- Adding a New Rule Category, page 6-14

Moving a Rule

License: Any

Proper access control rule order reduces the resources required to process network traffic, and prevents rule preemption.

The following procedure explains how to move one or more rules at a time using the access control policy editor. You can also move individual access control rules using the rule editor; see Creating and Editing Access Control Rules, page 6-2.

To move a rule:

Step 1 In the access control policy editor for the policy that contains the rules you want to move, select the rules by clicking in a blank area for each rule. Use the Ctrl and Shift keys to select multiple rules.

The rules you selected are highlighted.

Step 2 Move the rules. You can cut and paste or drag and drop.

To cut and paste rules into a new location, right-click a selected rule and select **Cut**. Then, right-click a blank area for a rule next to where you want to paste the cut rules and select **Paste above** or **Paste below**. Note that you cannot copy and paste access control rules between two different access control policies.

Step 3 Click **Store FirePOWER Changes** to save the policy.

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Adding a New Rule Category

License: Any

To help you organize access control rules, every access control policy has three system-provided rule categories: Administrator Rules, Standard Rules, and Root Rules. You cannot move, delete, or rename these categories, although you can create custom categories between the Standard Rules and Root Rules.

Adding custom categories allows you to further organize your rules without having to create additional policies. You can rename and delete categories that you add. You cannot move these categories, but you can move rules into, within, and out of them.

To add a new category:

Step 1 In the access control policy editor for the policy where you want to add a rule category, click Add Category.

Tip If your policy already contains rules, you can click a blank area in the row for an existing rule to set the position of the new category before you add it. You can also right-click an existing rule and select **Insert new category**.

The Add Category pop-up window appears.

Step 2 Type a unique category **Name**.

You can enter an alphanumeric name, including spaces and special printable characters, with up to 30 characters.

Step 3 You have the following choices:

- To position the new category immediately above an existing category, select **above Category** from the first **Insert** drop-down list, then select the category above which you want to position the rule from the second drop-down list.
- To position the new category rule below an existing rule, select **below rule** from the drop-down list, then enter an existing rule number. This option is valid only when at least one rule exists in the policy.
- To position the rule above an existing rule, select **above rule** from the drop-down list, then, enter an existing rule number. This option is valid only when at least one rule exists in the policy.

Step 4 Click **OK**.

Your category is added. You can click the edit icon ( ) next to a custom category to edit its name, or click the delete icon ( ) to delete the category. Rules in a category you delete are added to the category above.

Step 5 Click **Store FirePOWER Changes** to save the policy.



Controlling Traffic with Network-Based Rules

Access control rules within *access control policies* exert granular control over network traffic logging and handling. Network-based conditions allow you to manage which traffic can traverse your network, using one or more of the following criteria:

- source and destination security zones
- source and destination IP addresses or geographical locations
- source and destination port, which also includes transport layer protocol and ICMP code options

You can combine network-based conditions with each other and with other types of conditions to create an access control rule. These access control rules can be simple or complex, matching and inspecting traffic using multiple conditions. For detailed information on access control rules, see Tuning Traffic Flow Using Access Control Rules, page 6-1.

Note Security Intelligence-based traffic filtering, and some decoding and preprocessing occur **before** network traffic is evaluated by access control rules. You can also configure the *SSL inspection* feature to block or decrypt encrypted traffic before access control rules evaluate it.

Table 7-1 License Requirements for Network-Based Access Control Rules

Requirement	Geolocation Control	All Other Network-Based Control
license	Any	Any

For information on building network-based access control rules, see:

- Controlling Traffic by Security Zone, page 7-1
- Controlling Traffic by Network or Geographical Location, page 7-3
- Controlling Traffic by Port and ICMP Codes, page 7-5

Controlling Traffic by Security Zone

License: Any

Zone conditions in access control rules allow you to control traffic by its source and destination security zones. A *security zone* is a grouping of one or more interfaces.

As a simple example, you could create two zones: Internal and External, and assign the first pair of interfaces on the device to those zones. Hosts connected to the network on the Internal side represent your protected assets.

To extend this scenario, you could deploy additional identically configured devices to protect similar resources in several different locations. Each of these devices protects the assets in its Internal security zone.

Tip You are not required to group all internal (or external) interfaces into a single zone. Choose the grouping that makes sense for your deployment and security policies. For more information on creating zones, see Working with Security Zones, page 2-32.

In this deployment, you may decide that although you want these hosts to have unrestricted access to the Internet, you nevertheless want to protect them by inspecting incoming traffic for intrusions and malware.

To accomplish this using access control, configure an access control rule with a zone condition where the **Destination Zone** is set to **Internal**. This simple access control rule matches traffic that leaves the device from any interface in the Internal zone.

To ensure that the system inspects matching traffic for intrusions and malware, choose a rule action of **Allow**, then associate this rule with an intrusion and a file policy. For more information, see Using Rule Actions to Determine Traffic Handling and Inspection, page 6-7 and Controlling Traffic Using Intrusion and File Policies, page 11-1.

If you want to build a more complex rule, you can add a maximum of 50 zones to each of the **Source Zones** and **Destination Zones** in a single zone condition:

- To match traffic *leaving* the device from an interface in the zone, add that zone to the **Destination Zones**.

  Because devices deployed passively do not transmit traffic, you cannot use a zone comprised of passive interfaces in a **Destination Zone** condition.
- To match traffic *entering* the device from an interface in the zone, add that zone to the **Source Zones**.
- If you add both source and destination zone conditions to a rule, matching traffic must originate from one of the specified source zones **and** egress through one of the destination zones.

When building a zone condition, warning icons indicate invalid configurations. For details, Troubleshooting Access Control Policies and Rules, page 4-12.

To control traffic by zone:

Step 1 In the access control policy where you want to control traffic by zone, create a new access control rule or edit an existing rule.

For detailed instructions, see Creating and Editing Access Control Rules, page 6-2.

Step 2 In the rule editor, select the Zones tab.

The Zones tab appears.

Step 3 Find and select the zones you want to add from the **Available Zones**.

To search for zones to add, click the **Search by name** prompt above the **Available Zones** list, then type a zone name. The list updates as you type to display matching zones.

Click to select a zone. To select multiple zones, use the Shift and Ctrl keys, or right-click and then select **Select All**.

Step 4 Click **Add to Source** or **Add to Destination** to add the selected zones to the appropriate list.

You can also drag and drop selected zones.

Step 5 Save or continue editing the rule.

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Controlling Traffic by Network or Geographical Location

License: feature dependent

Network conditions in access control rules allow you to control traffic by its source and destination IP address. You can either:

- explicitly specify the source and destination IP addresses for the traffic you want to control, or
- use the geolocation feature, which associates IP addresses with geographical locations, to control traffic based on its source or destination country or continent

When you build a network-based access control rule condition, you can manually specify IP address and geographical locations. Alternately, you can configure network conditions with network and geolocation *objects*, which are reusable and associate a name with one or more IP addresses, address blocks, countries, continents, and so on.

Tip After you create a network or geolocation object, you can use it not only to build access control rules, but also to represent IP addresses in various other places in the system's module interface. For more information, see Managing Reusable Objects, page 2-1.

Note that if you want to write rules to control traffic by geographical location, to ensure you are using up-to-date geolocation data to filter your traffic, Cisco **strongly** recommends you regularly update the geolocation database (GeoDB) on your ASA FirePOWER module; see Updating the Geolocation Database, page 46-19.

Table 7-2 License Requirements for Network Conditions

Requirement	Geolocation Control	IP Address Control
license	Any	Any

You can add a maximum of 50 items to each of the **Source Networks** and **Destination Networks** in a single network condition, and you can mix network and geolocation-based configurations:

- To match traffic *from* an IP address or geographical location, configure the **Source Networks**.
- To match traffic *to* an IP address or geographical location, configure the **Destination Networks**.

If you add both source and destination network conditions to a rule, matching traffic must originate from one of the specified IP addresses **and** be destined for one of the destination IP addresses.

When building a network condition, warning icons indicate invalid configurations. For details, see Troubleshooting Access Control Policies and Rules, page 4-12.

Network conditions also allow you to handle proxied traffic based on the originating client. Use a source network condition to specify proxy servers, then add an original client constraint to specify original client IP addresses. The systems uses a packet's X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header field to determine original client IP.

Traffic matches the rule if the proxy's IP addresss matches the rule's source network constraint, and the original client's IP address matches the rule's original client constraint. For example, to allow traffic from a specific original client address, but only if it uses a specific proxy, create three rules:

Rule 1: Blocks non-proxied traffic from a specific IP address (209.165.201.1)

Source Networks: 209.165.201.1

Original Client Networks: none/any

Action: Block

Rule 2: Allows proxied traffic from the same IP address, but only if the proxy server for that traffic is one you choose (209.165.200.225 or 209.165.200.238)

Source Networks: 209.165.200.225 and 209.165.200.238

Original Client Networks: 209.165.201.1

Action: Allow

Rule 3: Blocks proxied traffic from the same IP address if it uses any other proxy server.

Source Networks: any

Original Client Networks: 209.165.201.1

Action: Block

To control traffic by network or geographical location:

Step 1 In the access control policy where you want to control traffic by network, create a new access control rule or edit an existing rule; see Creating and Editing Access Control Rules, page 6-2.

Step 2 In the rule editor, select the Networks tab.

Step 3 Find and select the networks you want to add from the **Available Networks**, as follows:

- Click the Networks tab to display network objects and groups to add; click the Geolocation tab to display geolocation objects.
- To add a network object on the fly, which you can then add to the condition, click the add icon ( ) above the **Available Networks** list; see Working with Network Objects, page 2-3.
- To search for network or geolocation objects to add, select the appropriate tab, click the **Search by name or value** prompt above the **Available Networks** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.

To select an object, click it. To select multiple objects, use the Shift and Ctrl keys, or right-click and then select **Select All**.

Step 4 If you want to filter proxied traffic:

- Click the **Source** sub-tab to specify a source network constraint.
- Click the **Original Client** sub-tab to specify an original client network constraint. In proxied connections, the original client's IP address must match one of these networks to match the rule.

Step 5 Click **Add to Source**, **Add to Original Client**, or **Add to Destination** to add the selected objects to the appropriate list.

You can also drag and drop selected objects.

Step 6 Add any source or destination IP addresses or address blocks that you want to specify manually.

Click the **Enter an IP address** prompt below the **Source Networks** or **Destination Networks** list; then type an IP address or address block and click **Add**.

Step 7 Save or continue editing the rule.

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Controlling Traffic by Port and ICMP Codes

License: Any

Network conditions in access control rules allow you to control traffic by its source and destination port. In this context, "port" refers to one of the following:

- For TCP and UDP, you can control traffic based on the transport layer protocol. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP(6)/22.
- For ICMP and ICMPv6 (IPv6-ICMP), you can control traffic based on its Internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- You can control traffic using other protocols that do not use ports.

When you build a port-based access control rule condition, you can manually specify ports. Alternately, you can configure port conditions with port *objects*, which are reusable and associate a name with one or more ports.

Tip After you create a port object, you can use it not only to build access control rules, but also to represent ports in various other places in the system's module interface. You can create port objects either using the object manager or on-the-fly while you are configuring access control rules. For more information, see Working with Port Objects, page 2-9.

You can add a maximum of 50 items to each of the **Selected Source Ports** and **Selected Destination Ports** lists in a single network condition:

- To match traffic *from* a port, configure the **Selected Source Ports**.

  If you add only source ports to a condition, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as source port conditions in a single access control rule.
- To match traffic *to* a port, configure the **Selected Destination Ports**.

  If you add only destination ports to a condition, you can add ports that use different transport protocols.
- To match traffic both originating from specific **Selected Source Ports** *and* destined for specific **Selected Destination Ports**, configure both.

If you add both source and destination ports to a condition, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

Keep the following points in mind when building a port condition:

- When you add a destination ICMP port with the type set to 0 or a destination ICMPv6 port with the type set to 129, the access control rule only matches unsolicited echo replies. ICMP echo replies sent in response to ICMP echo requests are ignored. For a rule to match on any ICMP echo, use ICMP type 8 or ICMPv6 type 128.
- When you use the GRE (47) protocol as a destination port condition, you can only add other network-based conditions to the access control rule, that is, zone, and network conditions. You cannot save the rule if you add reputation or user-based conditions.

When building a port condition, warning icons indicate invalid configurations. For example, you can use the object manager to edit in-use port objects so that the rules that use those object groups become invalid. For details, see Troubleshooting Access Control Policies and Rules, page 4-12.

To control traffic by port:

Step 1 In the access control policy where you want to control traffic by port, create a new access control rule or edit an existing rule.

For detailed instructions, see Creating and Editing Access Control Rules, page 6-2.

Step 2 In the rule editor, select the Ports tab.

The Ports tab appears.

Step 3 Find and select the ports you want to add from the **Available Ports**, as follows:

- To add a port object on the fly, which you can then add to the condition, click the add icon ( ) above the Available Ports list; see Working with Port Objects, page 2-9.
- To search for port objects and groups to add, click the **Search by name or value** prompt above the **Available Ports** list, then type either the name of the object, or the value of a port in the object. The list updates as you type to display matching objects. For example, if you type 80, the ASA FirePOWER module displays the Cisco-provided HTTP port object.

  To select an object, click it. To select multiple objects, use the Shift and Ctrl keys, or right-click and then select **Select All**.

Step 4 Click **Add to Source** or **Add to Destination** to add the selected objects to the appropriate list.

You can also drag and drop selected objects.

Step 5 Add any source or destination ports that you want to specify manually.

- For source ports, select either **TCP** or **UDP** from the **Protocol** drop-down list under the **Selected Source Ports** list. Then, enter a **Port**. You can specify a single port with a value from 0 to 65535.
- For destination ports, select a protocol (including **All** for all protocols) from the **Protocol** drop down list under the **Selected Destination Ports** list. You can also type the number of an unassigned protocol that does not appear in the list.

  If you select **ICMP** or **IPv6-ICMP**, a pop-up window appears where you can select a type and a related code. For more information on ICMP types and codes, see http://www.iana.org/assignments/icmp-parameters/icmp-parameters.xml and http://www.iana.org/assignments/icmpv6-parameters/icmpv6-parameters.xml.

If you do not want to specify a protocol, or optionally if you specified TCP or UDP, enter a **Port**. You can specify a single port with a value from 0 to 65535.

Click **Add**. Note that the ASA FirePOWER module will not add a port to a rule condition that results in an invalid configuration.

Step 6 Save or continue editing the rule.

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.



Controlling Traffic with Reputation-Based Rules

Access control rules within *access control policies* exert granular control over network traffic logging and handling. Reputation-based conditions in access control rules allow you to manage which traffic can traverse your network, by contextualizing your network traffic and limiting it where appropriate. Access control rules govern the following types of reputation-based control:

- Application conditions allow you to perform *application control*, which controls application traffic based on not only individual applications, but also applications' basic characteristics: type, risk, business relevance, categories, and tags.
- URL conditions allow you to perform *URL filtering*, which controls web traffic based on individual websites, as well as websites' system-assigned category and reputation.

You can combine reputation-based conditions with each other and with other types of conditions to create an access control rule. These access control rules can be simple or complex, matching and inspecting traffic using multiple conditions. For detailed information on access control rules, see Tuning Traffic Flow Using Access Control Rules, page 6-1.

Note Security Intelligence-based traffic filtering, and some decoding and preprocessing occur **before** network traffic is evaluated by access control rules. You can also configure the *SSL inspection* feature to block or decrypt encrypted traffic before access control rules evaluate it.

Reputation-based access control requires the following licenses.

Table 8-1 License Requirements for Reputation-Based Access Control Rules

Requirement	Application Control	URL Filtering (cat. & rep.)	URL Filtering (manual)
license	Control	URL Filtering	Any

For information on adding reputation-based conditions to access control rules, see:

- Controlling Application Traffic, page 8-2
- Blocking URLs, page 8-7

The ASA FirePOWER module can perform other types of reputation-based control, but you do not configure these using access control rules. For more information, see:

- Blacklisting Using Security Intelligence IP Address Reputation, page 5-1 explains how to limit traffic based on the reputation of a connection's origin or destination as a first line of defense.
- Tuning Intrusion Prevention Performance, page 11-6 explains how to detect, track, store, analyze, and block the transmission of malware and other types of prohibited files.

Controlling Application Traffic

License: Control

When the ASA FirePOWER module analyzes IP traffic, it can identify and classify the commonly used applications on your network.

Understanding Application Control

Application conditions in access control rules allow you to perform this *application control*. Within a single access control rule, there are a few ways you can specify applications whose traffic you want to control:

- You can select individual applications, including custom applications.
- You can use system-provided *application filters*, which are named sets of applications organized according to the applications' basic characteristics: type, risk, business relevance, categories, and tags.
- You can create and use custom application filters, which group applications (including custom applications) in any way you choose.

Application filters allow you to quickly create application conditions for access control rules. They simplify policy creation and administration, and grant you assurance that the system will control web traffic as expected. For example, you could create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the session is blocked.

In addition, Cisco frequently updates and adds additional detectors via system and vulnerability database (VDB) updates. By using filters based on application characteristics, you can ensure that the system uses the most up-to-date detectors to monitor application traffic.

Building Application Conditions

For traffic to match an access control rule with an application condition, the traffic must match one of the filters or applications that you add to a **Selected Applications and Filters** list.

In a single application condition, you can add a maximum of 50 items to the **Selected Applications and Filters** list. Each of the following counts as an item:

- One or more filters from the **Application Filters** list, individually or in custom combination. This item represents set of applications, grouped by characteristic.
- A filter created by saving an application search in the **Available Applications** list. This item represents a set of applications, grouped by substring match.
- An individual application from the **Available Applications** list.

In the module interface, filters added to a condition are listed above and separately from individually added applications.

Note that when you deploy an access control policy, for each rule with an application condition, the system generates a list of unique applications to match. In other words, you may use overlapping filters and individually specified applications to ensure complete coverage.

Note For encrypted traffic, the system can identify and filter traffic using only the applications tagged **SSL Protocol**. Applications without this tag can only be detected in unencrypted or decrypted traffic.

For more information, see the following sections:

- Matching Traffic with Application Filters, page 8-3
- Matching Traffic from Individual Applications, page 8-4
- Adding an Application Condition to an Access Control Rule, page 8-5
- Limitations to Application Control, page 8-6

Matching Traffic with Application Filters

License: Control

When building an application condition in an access control rule, use the **Application Filters** list to create a set of applications, grouped by characteristic, whose traffic you want to match.

Note that the mechanism for filtering applications within an access control rule is the same as that for creating reusable, custom application filters using the object manager; see Working with Application Filters, page 2-11. You can also save many filters you create on-the-fly in access control rules as new, reusable filters. You cannot save a filter that includes another user-created filter because you cannot nest user-created filters.

Understanding How Filters Are Combined

When you select filters, singly or in combination, the **Available Applications** list updates to display only the applications that meet your criteria. You can select system-provided filters in combination, but not custom filters.

The system links multiple filters of the same filter type with an OR operation. For example, if you select the Medium and High filters under the Risks type, the resulting filter is:

```
Risk: Medium OR High
```

If the Medium filter contains 110 applications and the High filter contains 82 applications, the system displays all 192 applications in the **Available Applications** list.

The system links different types of filters with an AND operation. For example, if you select the Medium and High filters under the Risks type, and the Medium and High filters under the Business Relevance type, the resulting filter is:

```
Risk: Medium OR High
AND
Business Relevance: Medium OR High
```

In this case, the system displays only those applications that are included in both the Medium or High Risk type AND the Medium or High Business Relevance type.

Finding and Selecting Filters

To select filters, click the arrow next to a filter type to expand it, then select or clear the check box next to each filter whose applications you want to display or hide. You can also right-click a system-provided filter type (**Risks**, **Business Relevance**, **Types**, **Categories**, or **Tags**) and select **Check All** or **Uncheck All**.

To search for filters, click the **Search by name** prompt above the **Available Filters** list, then type a name. The list updates as you type to display matching filters.

After you are done selecting filters, use the **Available Applications** list to add those filters to the rule; see Matching Traffic from Individual Applications, page 8-4.

Matching Traffic from Individual Applications

License: Control

When building an application condition in an access control rule, use the **Available Applications** list to select the applications whose traffic you want to match.

Browsing the List of Applications

When you first start to build the condition the list is unconstrained, and displays every application the system detects, 100 at a time:

- To page through the applications, click the arrows underneath the list.
- To display a pop-up window with summary information about the application's characteristics, as well as Internet search links that you can follow, click the information icon ( ) next to an application.

Finding Applications to Match

To help you find the applications you want to match, you can constrain the **Available Applications** list in the following ways:

- To search for applications, click the **Search by name** prompt above the list, then type a name. The list updates as you type to display matching applications.
- To constrain the applications by applying a filter, use the **Application Filters** list (see Matching Traffic with Application Filters, page 8-3). The **Available Applications** list updates as you apply filters.

Once constrained, an **All apps matching the filter** option appears at the top of the **Available Applications** list. This option allows you to add all the applications in the constrained list to the **Selected Applications and Filters** list, all at once.

Note If you select one or more filters in the Application Filters list and also search the **Available Applications** list, your selections and the search-filtered **Available Applications** list are combined using an AND operation. That is, the **All apps matching the filter** condition includes all the individual conditions currently displayed in the **Available Applications** list as well as the search string entered above the **Available Applications** list.

Selecting Single Applications to Match in a Condition

After you find an application you want to match, click to select it. To select multiple applications, use the Shift and Ctrl keys, or right-click and select **Select All** to select all applications in the current constrained view.

In a single application condition, you can match a maximum of 50 applications by selecting them individually; to add more than 50 you must either create multiple access control rules or use filters to group applications.

Selecting All Applications Matching a Filter for a Condition

Once constrained by either searching or using the filters in the **Application Filters** list, the **All apps matching the filter** option appears at the top of the **Available Applications** list.

This option allows you to add the entire set of applications in the constrained **Available Applications** list to the **Selected Applications and Filters** list, at once. In contrast to adding applications individually, adding this set of applications counts as only one item against the maximum of 50, regardless of the number of individual applications that comprise it.

When you build an application condition this way, the name of the filter you add to the **Selected Applications and Filters** list is a concatenation of the filter types represented in the filter plus the names of up to three filters for each type. More than three filters of the same type are followed by an ellipsis (...). For example, the following filter name includes two filters under the Risks type and four under Business Relevance:

```
Risks: Medium, High Business Relevance: Low, Medium, High,...
```

Filter types that are not represented in a filter you add with **All apps matching the filter** are not included in the name of the filter you add. These filter types are set to *any*; that is, these filter types do not constrain the filter, so any value is allowed for these.

You can add multiple instances of **All apps matching the filter** to an application condition, with each instance counting as a separate item in the **Selected Applications and Filters** list. For example, you could add all high risk applications as one item, clear your selections, then add all low business relevance applications as another item. This application condition matches applications that are high risk OR have low business relevance.

Adding an Application Condition to an Access Control Rule

License: Control

For traffic to match an access control rule with an application condition, the traffic must match one of the filters or applications that you add to a **Selected Applications and Filters** list.

You can add a maximum of 50 items per condition, and filters added to a condition are listed above and separately from individually added applications. When building an application condition, warning icons indicate invalid configurations. For details, see Troubleshooting Access Control Policies and Rules, page 4-12.

To control application traffic:

Step 1 In the access control policy where you want to control traffic by application, create a new access control rule or edit an existing rule.

For detailed instructions, see Creating and Editing Access Control Rules, page 6-2.

Step 2 In the rule editor, click the **Applications** tab.

Step 3 Optionally, enable content restriction features by clicking the dimmed icons for Safe Search ( ) or YouTube EDU ( ) and setting related options; for additional configuration requirements, see Using Access Control Rules to Enforce Content Restriction, page 13-1.

In most cases, enabling content restriction populates the condition's **Selected Applications and Filters** list with the appropriate values. The system does not automatically populate the list if applications or filters related to content restriction are already present in the list when you enable content restriction.

Continue with the procedure to refine your application and filter selections, or skip to saving the rule.

Step 4 Optionally, use filters to constrain the list of applications displayed in the **Available Applications** list.

Select one or more filters in the **Application Filters** list. For more information, see Matching Traffic with Application Filters, page 8-3.

Step 5 Find and select the applications you want to add from the **Available Applications** list.

You can search for and select individual applications, or, when the list is constrained, **All apps matching the filter**. For more information, see Matching Traffic from Individual Applications, page 8-4.

Step 6 Click **Add to Rule** to add the selected applications to the **Selected Applications and Filters** list.

You can also drag and drop selected applications and filters. Filters appear under the heading *Filters*, and applications appear under the heading *Applications*.

Tip Before you add another filter to this application condition, click **Clear All Filters** to clear your existing selections.

Step 7 Optionally, click the add icon ( ) above the **Selected Applications and Filters** list to save a custom filter comprised of all the individual applications and filters currently in the list.

Use the object manager to manage this on-the-fly-created filter; see Working with Application Filters, page 2-11. Note that you cannot save a filter that includes another user-created filter; you cannot nest user-created filters.

Step 8 Save or continue editing the rule.

You must deploy the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Limitations to Application Control

License: Control

Keep the following points in mind when performing application control.

Speed of Application Identification

The system cannot perform application control before:

- a monitored connection is established between a client and server, and
- the system identifies the application in the session

This identification should occur within 3 to 5 packets, or after the server certificate exchange in the SSL handshake if the traffic is encrypted. If one of these first packets matches all other conditions in an access control rule containing an application condition but the identification is not complete, the access control policy allows the packet to pass. This behavior allows the connection to be established so that applications can be identified. For your convenience, affected rules are marked with an information icon ( ).

The allowed packets are inspected by the access control policy's *default* intrusion policy (not the *default action* intrusion policy nor the almost-matched rule's intrusion policy). For more information, see Setting the Default Intrusion Policy for Access Control, page 20-1.

After the system completes its identification, the system applies the access control rule action, as well as any associated intrusion and file policy, to the remaining session traffic that matches its application condition.

Handling Encrypted Traffic

The system can identify and filter unencrypted application traffic that becomes encrypted using StartTLS, such as SMTPS, POPS, FTPS, TelnetS, and IMAPS. In addition, it can identify certain encrypted applications based on the Server Name Indication in the TLS client hello message, or the server certificate subject distinguished name value.

These applications are tagged **SSL Protocol**. Applications without this tag can only be detected in unencrypted or decrypted traffic.

Handling Application Traffic Packets Without Payloads

The system applies the default policy action to packets that do not have a payload in a connection where an application is identified.

Handling Referred Traffic

To create a rule to act on traffic referred by a web server, such as advertisement traffic, add a condition for the referred application rather than the referring application.

Controlling Application Traffic That Uses Multiple Protocols (Skype)

The system can detect multiple types of Skype application traffic. When building an application condition to control Skype traffic, select the **Skype** tag from the **Application Filters** list rather than selecting individual applications. This ensures that the system can detect and control all Skype traffic the same way. For more information, see Matching Traffic with Application Filters, page 8-3.

Blocking URLs

License: feature dependent

URL conditions in access control rules allow you to limit the websites that users on your network can access. This feature is called *URL filtering*. There are two ways you can use access control to specify URLs you want to block (or, conversely, allow):

- With any license, you can manually specify individual URLs or groups of URLs to achieve granular, custom control over web traffic.
- With a URL Filtering license, you can also control access to websites based on the URL's general classification, or *category*, and risk level, or *reputation*. The system displays this category and reputation data in connection logs, intrusion events, and application details.

Note To see URL category and reputation information in events, you must create at least one access control rule with a URL condition.

When you block a website, you can either allow the user's browser its default behavior, or you can display a generic system-provided or custom page. You can also give users a chance to bypass a website block by clicking through a warning page.

Table 8-2 License Requirements for URL Filtering

Requirement	Category & Reputation-Based	Manual
license	URL Filtering	Any

For more information, see:

- Performing Reputation-Based URL Blocking, page 8-8
- Performing Manual URL Blocking, page 8-10
- Limitations to URL Detection and Blocking, page 8-11
- Allowing Users to Bypass URL Blocks, page 8-12
- Displaying a Custom Web Page for Blocked URLs, page 8-14

Performing Reputation-Based URL Blocking

License: URL Filtering

With a URL Filtering license, you can control your users' access to websites based on the category and reputation of requested URLs, which the ASA FirePOWER module obtains from the Cisco cloud:

- The URL *category* is a general classification for the URL. For example, ebay.com belongs to the **Auctions** category, and monster.com belongs to the **Job Search** category. A URL can belong to more than one category.
- The URL *reputation* represents how likely the URL is to be used for purposes that might be against your organization's security policy. A URL's risk can range from **High Risk** (level 1) to **Well Known** (level 5).

Note Before access control rules with category and reputation-based URL conditions can take effect, you **must** enable communications with the Cisco cloud. This allows the ASA FirePOWER module to retrieve URL data. For more information, see Enabling Cloud Communications, page 44-2.

Advantages to Reputation-Based URL Blocking

URL categories and reputations allow you to quickly create URL conditions for access control rules. For example, you could create an access control rule that identifies and blocks all **High Risk** URLs in the **Abused Drugs** category. If a user attempts to browse to any URL with that category and reputation combination, the session is blocked.

Using category and reputation data from the Cisco cloud also simplifies policy creation and administration. It grants you assurance that the system will control web traffic as expected. Finally, because the cloud is continually updated with new URLs, as well as new categories and risks for existing URLs, you can ensure that the system uses up-to-date information to filter requested URLs. Malicious sites that represent security threats such as malware, spam, botnets, and phishing may appear and disappear faster than you can update and deploy new policies.

Some examples include:

- If a rule blocks all gaming sites, as new domains get registered and classified as **Gaming**, the system can block those sites automatically.
- If a rule blocks all malware sites, and a blog page gets infected with malware, the cloud can recategorize the URL from **Blog** to **Malware** and the system can block that site.
- If a rule blocks high-risk social networking sites, and somebody posts a link on their profile page that contains links to malicious payloads, the cloud can change the reputation of that page from **Benign sites** to **High Risk** so the system can block it.

Note that if the cloud does not know the category or reputation of a URL, or if the ASA FirePOWER module cannot contact the cloud, that URL does **not** trigger access control rules with category or reputation-based URL conditions. You cannot assign categories or reputations to URLs manually.

Building URL Conditions

You can add a maximum of 50 items to the **Selected URLs** to match in a single URL condition. Each URL category, optionally qualified by reputation, counts as a single item. Note that you can also use literal URLs and URL objects in URL conditions, but you cannot qualify these items with a reputation. For more information, see Performing Manual URL Blocking, page 8-10.

Note that you cannot qualify a literal URL or URL object with a reputation.

When building a URL condition, warning icons indicate invalid configurations. For details, see Troubleshooting Access Control Policies and Rules, page 4-12.

To control traffic by requested URL using category and reputation data:

Step 1 In the access control policy where you want to control traffic by URL, create a new access control rule or edit an existing rule.

For detailed instructions, see Creating and Editing Access Control Rules, page 6-2.

Step 2 In the rule editor, select the URLs tab.

The URLs tab appears.

Step 3 Find and select the categories of URL you want to add from the **Categories and URLs** list. To match web traffic regardless of category, select **Any** category.

To search for categories to add, click the **Search by name or value** prompt above the **Categories and URLs** list, then type the category name. The list updates as you type to display matching categories.

To select a category, click it. To select multiple categories, use the Shift and Ctrl keys.

Tip Although you can right-click and **Select All** categories, adding all categories this way exceeds the 50-item maximum for an access control rule. Instead, use **Any**.

Step 4 Optionally, qualify your category selections by clicking a reputation level from the **Reputations** list. If you do not specify a reputation level, the system defaults to **Any**, meaning all levels.

You can only select one reputation level. When you choose a reputation level, the access control rule behaves differently depending on its purpose:

- If the rule blocks or monitors web access (the rule action is **Block**, **Block with reset**, **Interactive Block**, **Interactive Block with reset**, or **Monitor**) selecting a reputation level also selects all reputations more severe than that level. For example, if you configure a rule to block or monitor **Suspicious sites** (level 2), it also automatically blocks or monitors **High risk** (level 1) sites.
- If the rule allows web access, whether to trust or further inspect it (the rule action is **Allow** or **Trust**), selecting a reputation level also selects all reputations less severe than that level. For example, if you configure a rule to allow **Benign sites** (level 4), it also automatically allows **Well known** (level 5) sites.

If you change the rule action for a rule, the system automatically changes the reputation levels in URL conditions according to the above points.

Step 5 Click **Add to Rule** or drag and drop the selected items to add them to the **Selected URLs** list.

Step 6 Save or continue editing the rule.

You must deploy the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Performing Manual URL Blocking

License: Any

To supplement or selectively override URL filtering by category and reputation, you can control web traffic by manually specifying individual URLs or groups of URLs. This allows you to achieve granular, custom control over allowed and blocked web traffic. You can also perform this type of URL filtering without a special license.

To manually specify URLs to allow or block in an access control rule, you can type in a single literal URL. Or, you can configure URL conditions using URL objects, which are reusable and associate a name with a URL or IP address.

Tip After you create a URL object, you can use it not only to build access control rules, but also to represent URLs in various other places in the system's module interface. You can create these objects using the object manager; you can also create URL objects on-the-fly while you are configuring access control rules. For more information, see Working with URL Objects, page 2-10.

Manually Specifying URLs in URL Conditions

Although manual entry gives you precise control over allowed and blocked web traffic, you cannot qualify a manually specified URL with a reputation. Additionally, you must make sure that your rules do not have unintended consequences. To determine whether network traffic matches a URL condition, the system performs a simple substring match. If the value of a URL object or manually typed URL matches any part of a URL requested by a monitored host, the URL condition of the access control rule is satisfied.

Therefore, when manually specifying URLs in URL conditions, including in URL objects, carefully consider other traffic that might be affected. For example, if you allow all traffic to example.com, your users could browse to URLs including:

- http://example.com/
- http://example.com/newexample
- http://www.example.com/

As another example, consider a scenario where you want to explicitly block ign.com (a gaming site). However, substring matching means that blocking *ign.com* also blocks veris*ign.com*, which might not be your intent.

Manually Blocking Encrypted Web Traffic

URL conditions in access control rules:

- disregard the encryption protocol of web traffic (HTTP vs HTTPS)

  For example, access control rules treat traffic to http://example.com/ the same as traffic to https://example.com/. To configure an access control rule that matches only HTTP or HTTPS traffic, add an application condition to the rule. For more information, see Blocking URLs, page 8-7.
- match HTTPS traffic based on the subject common name in the public key certificate used to encrypt the traffic, and also disregard subdomains within the subject common name

  Do not include subdomain information when manually filtering HTTPS traffic.

When building a URL condition, warning icons indicate invalid configurations. For details, see Troubleshooting Access Control Policies and Rules, page 4-12.

To control web traffic by manually specifying URLs to allow or block:

Step 1 In the access control policy where you want to control traffic by URL, create a new access control rule or edit an existing rule.

For detailed instructions, see Creating and Editing Access Control Rules, page 6-2.

Step 2 In the rule editor, select the URLs tab.

The URLs tab appears.

Step 3 Find and select the URL objects and groups you want to add from the **Categories and URLs** list:

- To add a URL object on the fly, which you can then add to the condition, click the add icon ( ) above the **Categories and URLs** list; see Working with URL Objects, page 2-10.
- To search for URL objects and groups to add, click the **Search by name or value** prompt above the **Categories and URLs** list, then type either the name of the object, or the value of a URL or IP address in the object. The list updates as you type to display matching objects.

To select an object, click it. To select multiple objects, use the Shift and Ctrl keys. Although you can right-click and **Select All** URL objects and categories, adding URLs this way exceeds the 50-item maximum for an access control rule.

Step 4 Click **Add to Rule** or to add the selected items to the **Selected URLs** list.

You can also drag and drop selected items.

Step 5 Add any literal URLs that you want to specify manually. You **cannot** use wildcards (*) in this field.

Click the **Enter URL** prompt below the **Selected URLs** list; then type a URL or IP address and click **Add**.

Step 6 Save or continue editing the rule.

You must deploy the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Limitations to URL Detection and Blocking

License: Any

Keep the following points in mind when performing URL detection and blocking.

Speed of URL Identification

The system cannot filter URLs before:

- a monitored connection is established between a client and server
- the system identifies the HTTP or HTTPS application in the session
- the system identifies the requested URL (for encrypted sessions, from either the client hello message or the server certificate)

This identification should occur within 3 to 5 packets, or after the server certificate exchange in the SSL handshake if the traffic is encrypted. If one of these first packets matches all other conditions in an access control rule containing a URL condition but the identification is not complete, the access control policy allows the packet to pass. This behavior allows the connection to be established so that URLs can be identified. For your convenience, affected rules are marked with an information icon ( ).

The allowed packets are inspected by the access control policy's *default* intrusion policy (not the *default action* intrusion policy nor the almost-matched rule's intrusion policy). For more information, see Setting the Default Intrusion Policy for Access Control, page 20-1.

After the system completes its identification, the system applies the access control rule action, as well as any associated intrusion and file policy, to the remaining session traffic that matches its URL condition.

Handling Encrypted Web Traffic

When evaluating encrypted web traffic using access control rules with URL conditions, the system:

- disregards the encryption protocol; an access control rule matches both HTTPS and HTTP traffic if the rule has a URL condition but not an application condition that specifies the protocol
- matches HTTPS traffic based on the subject common name in the public key certificate used to encrypt the traffic, and disregards subdomains within the subject common name
- does not display an HTTP response page, even if you configured one

Search Query Parameters in URLs

The system does not use search query parameters in the URL to match URL conditions. For example, consider a scenario where you block all shopping traffic. In that case, using a web search to search for amazon.com is not blocked, but browsing to amazon.com is.

Allowing Users to Bypass URL Blocks

License: Any

When you block a user's HTTP web request using an access control rule, setting the rule action to **Interactive Block** or **Interactive Block with reset** gives that user a chance to bypass the block by clicking through a warning *HTTP response page*. You can display a generic system-provided response page or you can enter custom HTML.

By default, the system allows users to bypass blocks for 10 minutes (600 seconds) without displaying the warning page on subsequent visits. You can set the duration to as long as a year, or you can force the user to bypass the block every time.

If the user does not bypass the block, matching traffic is denied without further inspection; you can also reset the connection. On the other hand, if a user bypasses the block, the system allows the traffic. Allowing this traffic means that you can continue to inspect unencrypted payloads for intrusions, malware and prohibited files. Note that users may have to refresh after bypassing the block to load page elements that did not load.

Note that you configure the interactive HTTP response page separately from the response page you configure for Block rules. For example, you could display the system-provided page to users whose sessions are blocked without interaction, but a custom page to users who can click to continue. For more information, see Displaying a Custom Web Page for Blocked URLs, page 8-14.

If you block web traffic decrypted by the SSL inspection feature, the system encrypts the response page and sends it at the end of the reencrypted SSL stream.



Tip To quickly disable interactive blocking for all rules in an access control policy, display neither the system-provided page nor a custom page. This causes the system to block all connections that match an Interactive Block rule without interaction.

To allow users to bypass a website block:

Step 1 Create an access control rule that matches web traffic with a URL condition.

See Performing Reputation-Based URL Blocking, page 8-8 and Performing Manual URL Blocking, page 8-10.

Step 2 Make sure the access control rule action is **Interactive Block** or **Interactive Block with reset.**

See Using Rule Actions to Determine Traffic Handling and Inspection, page 6-7.

Step 3 Assume users will bypass the block and choose inspection and logging options for the rule accordingly. As with Allow rules:

- You can associate either type of Interactive Block rule with a file and intrusion policy. For more information, see Controlling Traffic Using Intrusion and File Policies, page 11-1.
- Logging options for interactively blocked traffic are identical to those in allowed traffic, but keep in mind that if a user does not bypass the interactive block, the system can log only beginning-of-connection events.

  Note that when the system initially warns the user, it marks any logged beginning-of-connection event with the Interactive Block or Interactive Block with reset action. If the user bypasses the block, additional connection events logged for the session have an action of Allow. For more information, see Logging Connections Based on Access Control Handling, page 36-9.

Step 4 Optionally, set the amount of time that elapses after a user bypasses a block before the system displays the warning page again.

See Setting the User Bypass Timeout for a Blocked Website, page 8-13.

Step 5 Optionally, create and use a custom page to display to allow users to bypass a block.

See Displaying a Custom Web Page for Blocked URLs, page 8-14.

Setting the User Bypass Timeout for a Blocked Website

License: Any

By default, the system allows a user to bypass interactive blocks for 10 minutes (600 seconds) without displaying the warning page on subsequent visits. You can set the duration to as long as a year, or to zero to force the user to bypass the block every time. This limit applies to every Interactive Block rule in the policy. You cannot set the limit per rule.

To customize the length of time before a user bypass expires:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to configure.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

Advanced settings for the access control policy appear.

Step 4 Click the edit icon ( ) next to General Settings.

The General Settings pop-up window appears.

Step 5 In the **Allow an Interactive Block to bypass blocking for (seconds)** field, type the number of seconds that must elapse before the user bypass expires.

You can specify any number of seconds from zero to 31536000 (one year). Specifying zero forces your users to bypass the block every time.

Step 6 Click **OK**.

Advanced settings for the access control policy appear.

Step 7 Click **Store ASA FirePOWER Changes**.

You must deploy the access control policy for your changes to take effect. For more information, see Deploying Configuration Changes, page 4-11.

Displaying a Custom Web Page for Blocked URLs

License: Any

When the system blocks a user's HTTP web request, what the user sees in a browser depends on how you block the session, using the access control rule's action. You should select:

- **Block** or **Block with reset** to deny the connection. A blocked session times out; the system resets Block with reset connections. However, for both blocking actions, you can override the default browser or server page with a custom page that explains that the connection was denied. The system calls this custom page an *HTTP response page*.
- **Interactive Block** or **Interactive Block with reset** if you want to display an *interactive HTTP response page* that warns users, but also allows them to click a button to continue or refresh the page to load the originally requested site. Users may have to refresh after bypassing the response page to load page elements that did not load.

You can either display a generic system-provided response page, or you can enter custom HTML. When you enter custom text, a counter shows how many characters you have used.

In each access control policy, you configure the interactive HTTP response page separately from the response page you use to block traffic without interaction, that is, using a Block rule. For example, you could display the system-provided page to users whose sessions are blocked without interaction, but a custom page to users who can click to continue.

Reliable display of HTTP response pages to your users depends on your network configuration, traffic loads, and size of the page. If you build a custom response page, keep in mind that a smaller page is more likely to display successfully.

To configure HTTP response pages:

Step 1 Edit the access control policy monitoring your web traffic; see Editing Access Control Policies, page 4-7.

Step 2 Click the HTTP Responses tab.

Step 3 For the **Block Response Page** and the **Interactive Block Response Page**, choose responses from the drop-down lists. For each page, you have the following choices:

- **System-provided**—Displays a generic response. Click the view icon ( ) to view the code for this page.
- **Custom**—Create a custom response page.

  A pop-up window appears, prepopulated with system-provided code that you can replace or modify. When you are done, save your changes. Note that you can edit a custom page by clicking the edit icon ( ).
- **None**—Disables the response page and blocks sessions without interaction or explanation. Note that selecting this option for interactively blocked sessions prevents users from clicking to continue; the session is blocked without interaction.

Step 4 Click **Store ASA FirePOWER Changes**.

You must redeploy the configuration for your changes to take effect. For more information, see Deploying Configuration Changes, page 4-11.



Access Control Rules: Realms and Users

The following topics describe how to control user traffic on your network:

- Realm, User, User Group, and ISE Attribute Access Control Rule Conditions, page 9-1
- Troubleshooting Issues with User Access Control Rules, page 9-2
- Adding a Realm, User, or User Group Condition to an Access Control Rule, page 9-3
- Configuring ISE Attribute Conditions, page 9-3

Realm, User, User Group, and ISE Attribute Access Control Rule Conditions

License: Control

Before you can perform user control (create access control rule conditions based on entire realms, individual users, user groups, or ISE attributes), you must:

- Configure a realm for each Microsoft Active Directory or LDAP server you want to monitor. If you enable user download for the realm, the Firepower Management Center regularly and automatically queries the server to download metadata for newly or already-reported authoritative users and user groups.

Note Configuring a realm is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions.

- Create an identity policy to associate the realm with an authentication method.
- Configure one or more User Agents or ISE/ISE-PIC devices, or captive portal. If you want to use an ISE attribute condition, you must configure ISE.

  User Agents, ISE/ISE-PIC, and captive portal collect authoritative user data that can be used for user control in access control rule conditions. The identity sources monitor specified users as they log in or out of hosts or authenticate using LDAP or AD credentials.

Note If you configure a User Agent or ISE/ISE-PIC device to monitor a large number of user groups, or if you have a very large number of users mapped to hosts on your network, the system may drop user mappings based on groups, due to your Firepower Management Center user limit. As a result, access control rules with realm, user, or user group conditions may not fire as expected.

You can add a maximum of 50 realms, users, and groups to the Selected Users in a single user condition. Conditions with user groups match traffic to or from any of the group's members, including members of any sub-groups, with the exception of individually excluded users and members of excluded sub-groups.

Including a user group automatically includes all of that group's members, including members of any secondary groups. However, if you want to use the secondary group in access control rules, you must explicitly include the secondary group.

Note Hardware-based fast-path rules, Security Intelligence-based traffic filtering, SSL inspection, user identification, and some decoding and preprocessing occur before access control rules evaluate network traffic.

Troubleshooting Issues with User Access Control Rules

License: Control

If you notice unexpected user access control rule behavior, consider tuning your rule, identity source, or realm configurations.

Access control rules targeting realms, users, or user groups are not firing

If you configure a User Agent or ISE/ISE-PIC device to monitor a large number of user groups, or if you have a very large number of users mapped to hosts on your network, the system may drop user records due to your Firepower Management Center user limit. As a result, access control rules with realm or user conditions may not fire as expected.

Access control rules targeting user groups or users within user groups are not firing as expected

If you configure an access control rule with a user group condition, your LDAP or Active Directory server must have user groups configured. The Firepower Management Center cannot perform user group control if the server organizes the users in basic object hierarchy.

Access control rules targeting users in secondary groups are not firing as expected

If you configure an access control rule with a user group condition that includes or excludes users who are members of a secondary group on your Active Directory server, your server may be limiting the number of users it reports.

By default, Active Directory servers limit the number of users they report from secondary groups. You must customize this limit so that all of the users in your secondary groups are reported to the Firepower Management Center and eligible for use in access control rules with user conditions.

Access control rules are not matching users when seen for the first time

After the system detects activity from a previously-unseen user, the system retrieves information from the server. Until the system successfully retrieves this information, activity seen by this user is not handled by matching access control rules. Instead, the user session is handled by the next access control rule it matches (or the access control policy default action).

For example, this may explain when:

- Users who are members of user groups are not matching access control rules with user group conditions.
- Users who were reported by ISE/ISE-PIC or the User Agent are not matching access control rules, when the server used for user data retrieval is an Active Directory server.

Note that this may also cause the system to delay the display of user data in event views and analysis tools.

Adding a Realm, User, or User Group Condition to an Access Control Rule

License: Control

Before You Begin

- Configure one or more authoritative user identity sources as described in User Identity Sources, page 33-1.
- Configure a realm as described in Creating a Realm, page 32-4. A user download (automatic or on-demand) must be performed before you can configure realm, user, or user group conditions in an access control rule.

Step 1 In the access control rule editor, select the **Users** tab.

Step 2 Search by name or value above the **Available Realms** list and select a realm.

Step 3 Search by name or value above the **Available Users** list and select a user or group.

Step 4 Click **Add to Rule**, or drag and drop.

Step 5 Save or continue editing the rule.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Configuring ISE Attribute Conditions

License: Control

Before You Begin

- Configure a realm as described in Creating a Realm, page 32-4. A user download (automatic or on-demand) must be performed before you can configure an ISE attribute condition in an access control rule.

Note Configuring a realm is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions.

- Configure ISE as described in Configuring an ISE/ISE-PIC Connection, page 33-6.

Note The ISE-PIC identity source does not provide ISE attribute data. You must configure ISE.

Step 1 In the access control rule editor, click the **SGT/ISE Attributes** tab.

Step 2 Search by name or value above the **Available Attributes** list and choose an attribute.

Step 3 Search by name or value above the **Available Metadata** list and choose metadata.

Step 4 Click **Add to Rule**, or drag and drop.

Step 5 Constrain the rule with an IP address in the **Add a Location IP Address** field.

Note You can use ISE-assigned Security Group Tags (SGTs) to constrain ISE attribute conditions. To use custom SGTs in access control rules, see ISE SGT v. Custom SGT Rule Conditions, page 10-1.

Step 6 Save or continue editing the rule.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.



Access Control Rules: Custom Security Group Tags

The Security Group Tag (SGT) specifies the privileges of a traffic source within a trusted network. Security Group Access (a feature of both Cisco TrustSec and Cisco ISE) automatically generates the SGT when a user adds a security group in TrustSec or ISE. SGA then applies the SGT attribute as packets enter the network. You can use SGTs for access control by configuring ISE as an identity source or creating custom SGT objects.

Custom SGT conditions allow you to configure access control rules based on custom SGT objects. You manually add custom SGT objects to the Firepower System, rather than obtaining SGTs via ISE.

You can only use custom SGT conditions if you disable ISE/ISE-PIC as an identity source.

The following topics describe how to use SGT conditions in access control rules:

- ISE SGT v. Custom SGT Rule Conditions, page 10-1
- Automatic Transition from Custom SGT to ISE SGT Rule Conditions, page 10-2
- Configuring Custom SGT Conditions, page 10-2
- Troubleshooting Custom SGT Conditions, page 10-3

ISE SGT v. Custom SGT Rule Conditions

You can use SGTs for access control by either configuring ISE as an identity source (*ISE SGT*) or creating custom SGT objects (*custom SGT*). The system handles ISE SGT and custom SGT rule conditions differently:

ISE SGT: ISE connection configured

You can use ISE SGTs as ISE attribute conditions in access control rules. When you choose **Security Group Tag** from the **Available Attributes** list in the **SGT/ISE Attributes** tab, the system populates the **Available Metadata** list by querying ISE for available tags. The presence or absence of an SGT attribute in a packet determines the system's response:

- If an SGT attribute is present in the packet, the system extracts that value and compares it to ISE SGT conditions in access control rules.
- If the SGT attribute is absent from the packet, the system determines whether the SGT associated with the packet's source IP address is known in ISE and compares the SGT to the ISE SGT conditions in access control rules.

Custom SGT: No ISE connection configured

You can create custom SGT objects and use them as conditions in access control rules. When you choose **Security Group Tag** from the **Available Attributes** list in the **SGT/ISE Attributes** tab, the system populates the **Available Metadata** list with any SGT objects you have added. The presence or absence of an SGT attribute in a packet determines the system's response:

- If an SGT attribute is present in the packet, the system extracts that value and compares it to custom SGT conditions in access control rules.
- If the SGT attribute is absent from the packet, the system does not match the packet to custom SGT conditions in access control rules.

Automatic Transition from Custom SGT to ISE SGT Rule Conditions

If you create access control rules using custom SGT objects as conditions, then later configure ISE/ISE-PIC as an identity source, the system:

- Disables the **Security Group Tag** object option in the Object Manager. You cannot add new SGT objects, edit existing SGT objects, or add SGT objects as new conditions unless you disable the ISE/ISE-PIC connection.
- Retains existing SGT objects. You cannot modify these existing objects. You can view them only in the context of the existing access control rules that use them as conditions.
- Retains existing access control rules with custom SGT conditions. Because custom SGT objects can only be updated via manual editing, Cisco recommends that you delete or disable these rules. Instead, create rules using SGTs as ISE attribute conditions. The system automatically queries ISE to update SGT metadata for ISE attribute conditions, but you can only update custom SGT objects via manual editing.

Configuring Custom SGT Conditions

License: Any

To configure a custom Security Group Tag (SGT) condition:

Step 1 In the access control rule editor, click the **SGT/ISE Attributes** tab.

Step 2 Choose **Security Group Tag** from the **Available Attributes** list.

Step 3 In the **Available Metadata** list, find and choose a custom SGT object.

If you choose , the rule matches all traffic with an SGT attribute. For example, you might choose this value if you want the rule to block traffic from hosts that are not configured for TrustSec.

Step 4 Click **Add to Rule**, or drag and drop.

Step 5 Save or continue editing the rule.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Troubleshooting Custom SGT Conditions

If you notice unexpected rule behavior, consider tuning your custom SGT object configuration.

Security Group Tag objects unavailable

Custom SGT objects are only available if you do not configure ISE/ISE-PIC as an identity source. For more information, see Automatic Transition from Custom SGT to ISE SGT Rule Conditions, page 10-2.



Controlling Traffic Using Intrusion and File Policies

Intrusion and file policies work together, as the last line of defense before traffic is allowed to its destination:

- **Intrusion policies** govern the system's intrusion prevention capabilities; see Understanding Network Analysis and Intrusion Policies, page 18-1.
- **File policies** govern the system's network-based file control and advanced malware protection (AMP) capabilities; see Understanding and Creating File Policies, page 35-4.

Security Intelligence-based traffic filtering (blacklisting), SSL inspection-based decisions, and traffic decoding and preprocessing occur **before** network traffic is examined for intrusions, prohibited files, and malware. Access control rules and the access control default action determine which traffic is inspected by intrusion and file policies.

By associating an intrusion or file policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy, a file policy, or both.

Intrusion prevention and AMP require that you enable specific licensed capabilities as described in the following table.

Table 11-1 License Requirements for Intrusion and File Inspection

Feature	Description	License
intrusion prevention	detect and optionally block intrusions and exploits	Protection
file control	detect and optionally block the transmission of file types	Protection
advanced malware protection (AMP)	detect, track, and optionally block the transmission of malware	Malware

For more information on inspecting traffic for intrusions, prohibited files, and malware, see:

- Inspecting Allowed Traffic For Intrusions and Malware, page 11-2
- Tuning Intrusion Prevention Performance, page 11-6
- Tuning File and Malware Inspection Performance and Storage, page 11-16

Inspecting Allowed Traffic For Intrusions and Malware

License: Protection or Malware

Intrusion and file policies govern the system's intrusion prevention, file control, and AMP capabilities as a last line of defense before traffic is allowed to its destination. Security Intelligence-based traffic filtering, SSL inspection decisions (including decryption), decoding and preprocessing, and access control rule selection occur **before** intrusion and file inspection.

By associating an intrusion or file policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy, a file policy, or both. Access control rule conditions can be simple or complex; you can control traffic by security zone, network or geographical location, port, application, requested URL, and user.

The system matches traffic to access control rules in the order you specify. In most cases, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic. An access control rule's *action* determines how the system handles matching traffic. You can monitor, trust, block, or allow (with or without further inspection) matching traffic; see Using Rule Actions to Determine Traffic Handling and Inspection, page 6-7.

Note that an Interactive Block rule has the same inspection options as an Allow rule. This is so you can inspect traffic for malicious content when a user bypasses a blocked website by clicking through a warning page. For more information, see Interactive Blocking Actions: Allowing Users to Bypass Website Blocks, page 6-9.

Traffic that does not match any of the non-Monitor access control rules in a policy is handled by the default action. Note that the system can inspect traffic allowed by the default action for intrusions, but not prohibited files or malware. You **cannot** associate a file policy with the access control default action.

Note Sometimes, when a connection is analyzed by an access control policy, the system must process the first few packets in that connection, **allowing them to pass**, before it can decide which access control rule (if any) will handle the traffic. However, so these packets do not reach their destination uninspected, you can use an intrusion policy—called the default intrusion policy—to inspect them and generate intrusion events. For more information, see Setting the Default Intrusion Policy for Access Control, page 20-1.

For more information on the above scenario and instructions on associating file and intrusion policies with access control rules and the access control default action, see:

- Understanding File and Intrusion Inspection Order, page 11-2
- Configuring an Access Control Rule to Perform AMP or File Control, page 11-3
- Configuring an Access Control Rule to Perform Intrusion Prevention, page 11-4
- Setting Default Handling and Inspection for Network Traffic, page 4-4

Understanding File and Intrusion Inspection Order

License: Protection or Malware

Note Traffic allowed by an Intrusion Prevention default action can be inspected for intrusions, but cannot be inspected for prohibited files or malware. You **cannot** associate a file policy with the access control default action.

You do not have to perform both file and intrusion inspection in the same rule. For a connection matching an Allow or Interactive Block rule:

- without a file policy, traffic flow is determined by the intrusion policy
- without an intrusion policy, traffic flow is determined by the file policy

Tip The system does not perform any kind of inspection on trusted traffic.

For any single connection handled by an access control rule, file inspection occurs before intrusion inspection. That is, the system does not inspect files blocked by a file policy for intrusions. Within file inspection, simple blocking by type takes precedence over malware inspection and blocking.

Note Until a file is detected and blocked in a session, packets from the session may be subject to intrusion inspection.

For example, consider a scenario where you normally want to allow certain network traffic as defined in an access control rule. However, as a precaution, you want to block the download of executable files, examine downloaded PDFs for malware and block any instances you find, and perform intrusion inspection on the traffic.

You create an access control policy with a rule that matches the characteristics of the traffic you want to provisionally allow, and associate it with both an intrusion policy and a file policy. The file policy blocks the download of all executables, and also inspects and blocks PDFs containing malware:

- First, the system blocks the download of all executables, based on simple type matching specified in the file policy. Because they are immediately blocked, these files are subject to neither malware cloud lookup nor intrusion inspection.
- Next, the system performs malware cloud lookups for PDFs downloaded to a host on your network. Any PDFs with a malware file disposition are blocked, and are not subject to intrusion inspection.
- Finally, the system uses the intrusion policy associated with the access control rule to inspect any remaining traffic, including files not blocked by the file policy.

Configuring an Access Control Rule to Perform AMP or File Control

License: Protection or Malware

An access control policy can have multiple access control rules associated with file policies. You can configure file inspection for any Allow or Interactive Block access control rule, which permits you to match different file and malware inspection profiles against different types of traffic on your network before it reaches its final destination.

When the system detects a prohibited file (including malware) according to the settings in the file policy, it automatically logs an event. If you do not want to log file or malware events, you can disable this logging on a per-access-control-rule basis. After you associate the file policy with the access control rule, clear the **Log Files** check box on the Logging tab of the access control rule editor. For more information, see Disabling File and Malware Event Logging for Allowed Connections, page 36-7.

The system also logs the end of the associated connection, regardless of the logging configuration of the invoking access control rule; see Connections Associated with File and Malware Events (Automatic), page 36-3.

To associate a file policy with an access control rule:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy where you want to configure AMP or file control using access control rules.

Step 3 Create a new rule or edit an existing rule; see Creating and Editing Access Control Rules, page 6-2.

The access control rule editor appears.

Step 4 Ensure the rule action is set to **Allow**, **Interactive Block**, or **Interactive Block with reset**.

Step 5 Select the Inspection tab.

The Inspection tab appears.

Step 6 Select a **File Policy** to inspect traffic that matches the access control rule, or select **None** to disable file inspection for matching traffic.

You can click the edit icon ( ) that appears to edit the policy; see Creating a File Policy, page 35-9.

Step 7 Click **Add** to save the rule.

Your rule is saved. You must save and apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Configuring an Access Control Rule to Perform Intrusion Prevention

License: Protection

An access control policy can have multiple access control rules associated with intrusion policies. You can configure intrusion inspection for any Allow or Interactive Block access control rule, which permits you to match different intrusion inspection profiles against different types of traffic on your network before it reaches its final destination.

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.

Tip Even if you use system-provided intrusion policies, Cisco **strongly** recommends you configure the system's intrusion variables to accurately reflect your network environment. At a minimum, modify default variables in the default set; see Optimizing Predefined Default Variables, page 2-13.

Although you can associate a different intrusion policy-variable set pair with each Allow and Interactive Block rule (as well as with the default action), you cannot apply an access control policy if the target devices have insufficient resources to perform inspection as configured. For more information, see Simplifying Rules to Improve Performance, page 4-13.

Understanding System-Provided and Custom Intrusion Policies

Cisco delivers several intrusion policies with the ASA FirePOWER module. By using system-provided intrusion policies, you can take advantage of the experience of the Cisco Vulnerability Research Team (VRT). For these policies, the VRT sets intrusion and preprocessor rule states, as well as provides the

initial configurations for advanced settings. You can use system-provided policies as-is, or you can use them as the base for custom policies. Building custom policies can improve the performance of the system in your environment and provide a focused view of the malicious traffic and policy violations occurring on your network.

In addition to custom policies that you create, the system provides two custom policies: Initial Inline Policy and Initial Passive Policy. These two intrusion policies use the Balanced Security and Connectivity intrusion policy as their base. The only difference between them is their **Drop When Inline** setting, which enables drop behavior in the inline policy and disables it in the passive policy. For more information, see Comparing System-Provided with Custom Policies, page 18-7.

Connection and Intrusion Event Logging

When an intrusion policy invoked by an access control rule detects an intrusion, it generates an intrusion event. The system also automatically logs the end of the connection where the intrusion occurred, regardless of the logging configuration of the access control rule; see Connections Associated with Intrusions (Automatic), page 36-2.

To associate an intrusion policy with an access control rule:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy where you want to configure intrusion inspection using access control rules.

Step 3 Create a new rule or edit an existing rule; see Creating and Editing Access Control Rules, page 6-2.

The access control rule editor appears.

Step 4 Ensure the rule action is set to **Allow**, **Interactive Block**, or **Interactive Block with reset**.

Step 5 Select the Inspection tab.

The Inspection tab appears.

Step 6 Select a system-provided or custom **Intrusion Policy**, or select **None** to disable intrusion inspection for traffic that matches the access control rule.

If you select a custom intrusion policy, you can click the edit icon ( ) that appears to edit the policy; see Editing Intrusion Policies, page 26-4.

Caution Do **not** select `Experimental Policy 1` unless instructed to by a Cisco representative. Cisco uses this policy for testing.

Step 7 Optionally, change the **Variable Set** associated with the intrusion policy.

You can click the edit icon ( ) that appears to edit the variable set; see Working with Variable Sets, page 2-13.

Step 8 Click **Save** to save the rule.

Your rule is saved. You must save and apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Tuning Intrusion Prevention Performance

License: Protection

Cisco provides several features for improving the performance of your system as it analyzes traffic for attempted intrusions. You configure these performance settings on a per-access-control-policy basis, and they apply to all intrusion policies invoked by that parent access control policy.

For more information, see:

- Limiting Pattern Matching for Intrusions, page 11-6 describes how you can specify the number of packets to allow in the event queue, and enable or disable inspection of packets that will be rebuilt into larger streams.
- Overriding Regular Expression Limits for Intrusion Rules, page 11-7 describes how you can override default match and recursion limits on Perl-compatible regular expressions (PCRE).
- Limiting Intrusion Events Generated Per Packet, page 11-8 describes how you can configure rule processing event queue settings.
- Configuring Packet and Intrusion Rule Latency Thresholds, page 11-9 describes how you can balance security with the need to maintain device latency at an acceptable level with packet and rule latency thresholding.
- Configuring Intrusion Performance Statistic Logging, page 11-15 describes how you can configure basic performance monitoring and reporting parameters.

Limiting Pattern Matching for Intrusions

License: Protection

You can specify the number of packets to allow in the event queue. You can also, before and after stream reassembly, enable or disable inspection of packets that will be rebuilt into larger streams.

To configure event queue settings:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy.**

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the Advanced tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Performance Settings**, then select the **Pattern Matching Limits** tab in the pop-up window that appears.

Step 5 You can modify the following options:

- Type a value for the maximum number of events to queue in the **Maximum Pattern States to Analyze Per Packet** field.
- To inspect packets which will be rebuilt into larger streams of data before and after stream reassembly, select **Disable Content Checks on Traffic Subject to Future Reassembly**. Inspection before and after reassembly requires more processing overhead and may decrease performance.
- To disable inspection of packets which will be rebuilt into larger streams of data before and after stream reassembly, clear **Disable Content Checks on Traffic Subject to Future Reassembly**. Disabling inspection decreases the processing overhead for inspection of stream inserts and may boost performance.

Step 6 Click **OK**.

You must save and apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Overriding Regular Expression Limits for Intrusion Rules

License: Protection

You can override default match and recursion limits on PCRE that are used in intrusion rules to examine packet payload content. See Searching for Content Using PCRE, page 30-35 for information on using the `pcre` keyword in intrusion rules. The default limits ensure a minimum level of performance. Overriding these limits could increase security, but could also significantly impact performance by permitting packet evaluation against inefficient regular expressions.

Caution Do not override default PCRE limits unless you are an experienced intrusion rule writer with knowledge of the impact of degenerative patterns.

The following table describes the options you can configure to override the default limits.

Table 11-2 Regular Expression Constraint Options

Option	Description
Match Limit State	Specifies whether to override **Match Limit**. You have the following options: • select **Default** to use the value configured for **Match Limit** • select **Unlimited** to permit an unlimited number of attempts • select **Custom** to specify either a limit of 1 or greater for **Match Limit**, or to specify 0 to completely disable PCRE match evaluations
Match Limit	Specifies the number of times to attempt to match a pattern defined in a PCRE regular expression.

Table 11-2 Regular Expression Constraint Options (continued)

Option	Description
Match Recursion Limit State	Specifies whether to override **Match Recursion Limit**. You have the following options: • select **Default** to use the value configured for **Match Recursion Limit** • select **Unlimited** to permit an unlimited number of recursions • select **Custom** to specify either a limit of 1 or greater for **Match Recursion Limit**, or to specify 0 to completely disable PCRE recursions Note that for **Match Recursion Limit** to be meaningful, it must be smaller than **Match Limit**.
Match Recursion Limit	Specifies the number of recursions when evaluating a PCRE regular expression against the packet payload.

To configure PCRE overrides:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy.**

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the Advanced tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Performance Settings**, then select the Regular Expression Limits tab in the pop-up window that appears.

Step 5 You can modify any of the options in the Regular Expression Constraint Options table.

Step 6 Click **OK**.

You must save and apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Limiting Intrusion Events Generated Per Packet

License: Protection

When the rules engine evaluates traffic against rules, it places the events generated for a given packet or packet stream in an event queue, then reports the top events in the queue to the user interface. You can elect to have the rules engine log more than one event per packet or packet stream when multiple events are generated. Logging these events allows you to collect information beyond the reported event. When configuring this option, you can specify how many events can be placed in the queue and how many are logged, and select the criteria for determining event order within the queue.

The following table describes the options you can configure to determine how many events are logged per packet or stream.

Table 11-3 Intrusion Event Logging Limits Options

Option	Description
Maximum Events Stored Per Packet	The maximum number of events that can be stored for a given packet or packet stream.
Maximum Events Logged Per Packet	The number of events logged for a given packet or packet stream. This cannot exceed the **Maximum Events Stored Per Packet** value.
Prioritize Event Logging By	The value used to determine event ordering within the event queue. The highest ordered event is reported through the user interface. You can select from: • `priority`, which orders events in the queue by the event priority. • `content_length`, which orders events by the longest identified content match. When events are ordered by content length, rule events always take precedence over decoder and preprocessor events.

To configure how many events are logged per packet or stream:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy.**

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the Advanced tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Performance Settings**, then select the **Intrusion Event Logging Limits** tab in the pop-up window that appears.

Step 5 You can modify any of the options in the Intrusion Event Logging Limits Options table.

Step 6 Click **OK**.

You must save and apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Configuring Packet and Intrusion Rule Latency Thresholds

License: Protection

You can balance security with the need to maintain device latency at an acceptable level with packet and rule latency thresholding. For more information, see:

- Understanding Packet Latency Thresholding, page 11-10
- Configuring Packet Latency Thresholding, page 11-11
- Understanding Rule Latency Thresholding, page 11-12
- Configuring Rule Latency Thresholding, page 11-14

Understanding Packet Latency Thresholding

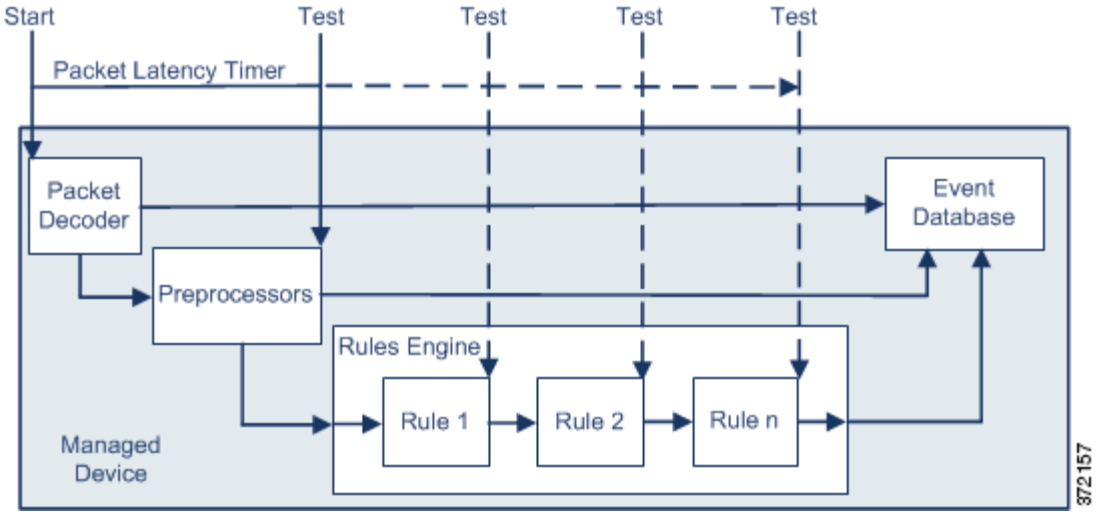
License: Protection

You can balance security with the need to maintain latency at an acceptable level by enabling packet latency thresholding. Packet latency thresholding measures the total elapsed time taken to process a packet by applicable decoders, preprocessors, and rules, and ceases inspection of the packet if the processing time exceeds a configurable threshold.

Packet latency thresholding measures elapsed time, not just processing time, in order to more accurately reflect the actual time required for the rule to process a packet. However, latency thresholding is a software-based latency implementation that does not enforce strict timing.

The trade-off for the performance and latency benefits derived from latency thresholding is that uninspected packets could contain attacks. However, packet latency thresholding gives you a tool you can use to balance security with connectivity.

A timer starts for each packet when decoder processing begins. Timing continues either until all processing ends for the packet or until the processing time exceeds the threshold at a timing test point.



As illustrated in the above figure, packet latency timing is tested at the following test points:

- after the completion of all decoder and preprocessor processing and before rule processing begins
- after processing by each rule

If the processing time exceeds the threshold at any test point, packet inspection ceases.

Tip Total packet processing time does not include routine TCP stream or IP fragment reassembly times.

Packet latency thresholding has no effect on events triggered by a decoder, preprocessor, or rule processing the packet. Any applicable decoder, preprocessor, or rule triggers normally until a packet is fully processed, or until packet processing ends because the latency threshold is exceeded, whichever comes first. If a drop rule detects an intrusion in an inline deployment, the drop rule triggers an event and the packet is dropped.

Note No packets are evaluated against rules after processing for that packet ceases because of a packet latency threshold violation. A rule that would have triggered an event cannot trigger that event, and for drop rules, cannot drop the packet.

For more information on drop rules, see Setting Rule States, page 27-19.

Packet latency thresholding can improve system performance in both passive and inline deployments, and can reduce latency in inline deployments, by stopping inspection of packets that require excessive processing time. These performance benefits might occur when, for example:

- for both passive and inline deployments, sequential inspection of a packet by multiple rules requires an excessive amount of time
- for inline deployments, a period of poor network performance, such as when someone downloads an extremely large file, slows packet processing

In a passive deployment, stopping the processing of packets might not contribute to restoring network performance because processing simply moves to the next packet.

Configuring Packet Latency Thresholding

License: Protection

The following table describes the options you can set to configure packet latency thresholding.

Table 11-4 Packet Latency Thresholding Options

Option	Description
Threshold (microseconds)	Specifies the time, in microseconds, when inspection of a packet ceases. See the Minimum Packet Latency Threshold Settings table for recommended minimum threshold settings.

You can enable rule 134:3 to generate an event when the system stops inspecting a packet because the packet latency threshold is exceeded. See Setting Rule States, page 27-19 for more information.

Many factors affect measurements of system performance and packet latency, such as CPU speed, data rate, packet size, and protocol type. For this reason, Cisco recommends that you use the threshold settings in the following table until your own calculations provide you with settings tailored to your network environment.

Table 11-5 Minimum Packet Latency Threshold Settings

For this data rate...	Set threshold microseconds to at least...
1 Gbps	100
100 Mbps	250
5 Mbps	1000

Determine the following when calculating your settings:

- average packets per second
- average microseconds per packet

Multiply the average microseconds per packet for your network by a significant safety factor to ensure that you do not unnecessarily discontinue packet inspections.

For example, the Minimum Packet Latency Threshold Settings table recommends a minimum packet latency threshold of 100 microseconds in a one gigabit environment. This minimum recommendation is based on test data showing an average of 250,000 packets per second, which is 0.25 packets per microsecond, or 4 microseconds per packet. Multiplying by a factor of twenty-five results in a recommended minimum threshold of 100 microseconds.

To configure packet latency thresholding:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy.**

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Latency-Based Performance Settings**, then select the Packet Handling tab in the pop-up window that appears.

Step 5 See the Minimum Packet Latency Threshold Settings table for recommended minimum **Threshold** settings.

Step 6 Click **OK**.

You must save and apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Understanding Rule Latency Thresholding

License: Protection

You can balance security with the need to maintain latency at an acceptable level by enabling rule latency thresholding. Rule latency thresholding measures the elapsed time each rule takes to process an individual packet, suspends the violating rule along with a group of related rules for a specified time if the processing time exceeds the rule latency threshold a configurable consecutive number of times, and restores the rules when the suspension expires.

Rule latency thresholding measures elapsed time, not just processing time, in order to more accurately reflect the actual time required for the rule to process a packet. However, latency thresholding is a software-based latency implementation that does not enforce strict timing.

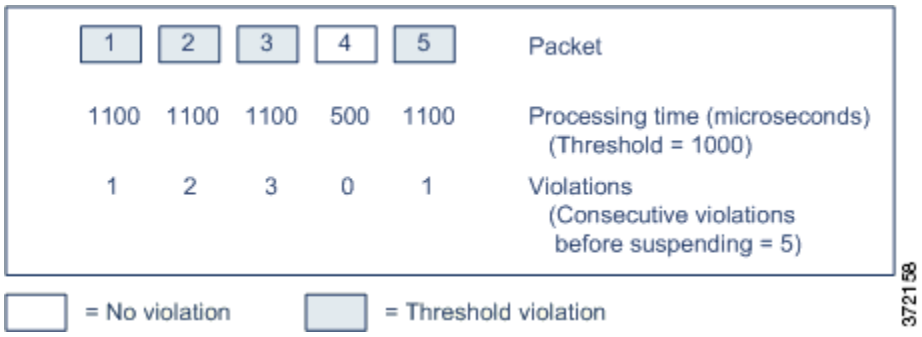
The trade-off for the performance and latency benefits derived from latency thresholding is that uninspected packets could contain attacks. However, rule latency thresholding gives you a tool you can use to balance security with connectivity.

A timer measures the processing time each time a packet is processed against a group of rules. Any time the rule processing time exceeds a specified rule latency threshold, the system increments a counter. If the number of consecutive threshold violations reaches a specified number, the system takes the following actions:

- suspends the rules for the specified period
- triggers an event indicating the rules have been suspended
- re-enables the rules when the suspension expires
- triggers an event indicating the rules have been re-enabled

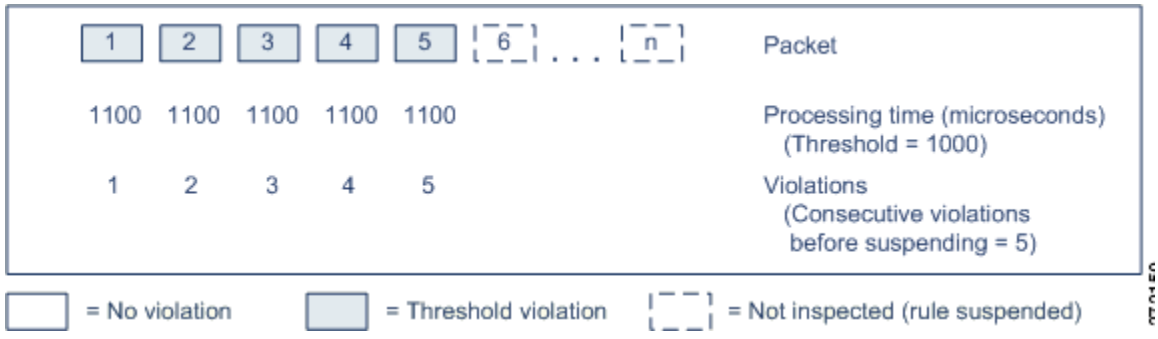
The system zeroes the counter when the group of rules has been suspended, or when rule violations are not consecutive. Permitting some consecutive violations before suspending rules lets you ignore occasional rule violations that might have negligible impact on performance and focus instead on the more significant impact of rules that repeatedly exceed the rule latency threshold.

The following example shows five consecutive rule processing times that do not result in rule suspension.



In the above example, the time required to process each of the first three packets violates the rule latency threshold of 1000 microseconds, and the violations counter increments with each violation. Processing of the fourth packet does not violate the threshold, and the violations counter resets to zero. The fifth packet violates the threshold and the violations counter restarts at one.

The following example shows five consecutive rule processing times that do result in rule suspension.



In the second example, the time required to process each of the five packets violates the rule latency threshold of 1000 microseconds. The group of rules is suspended because the rule processing time of 1100 microseconds for each packet violates the threshold of 1000 microseconds for the specified five

consecutive violations. Any subsequent packets, represented in the figure as packets 6 through n, are not examined against suspended rules until the suspension expires. If more packets occur after the rules are re-enabled, the violations counter begins again at zero.

Rule latency thresholding has no effect on intrusion events triggered by the rules processing the packet. A rule triggers an event for any intrusion detected in the packet, regardless of whether the rule processing time exceeds the threshold. If the rule detecting the intrusion is a drop rule in an inline deployment, the packet is dropped. When a drop rule detects an intrusion in a packet that results in the rule being suspended, the drop rule triggers an intrusion event, the packet is dropped, and that rule and all related rules are suspended. For more information on drop rules, see Setting Rule States, page 27-19.

Note Packets are not evaluated against suspended rules. A suspended rule that would have triggered an event cannot trigger that event and, for drop rules, cannot drop the packet.

Rule latency thresholding can improve system performance in both passive and inline deployments, and can reduce latency in inline deployments, by suspending rules that take the most time to process packets. Packets are not evaluated again against suspended rules until a configurable time expires, giving the overloaded device time to recover. These performance benefits might occur when, for example:

- hastily written, largely untested rules require an excessive amount of processing time
- a period of poor network performance, such as when someone downloads an extremely large file, causes slow packet inspection

Configuring Rule Latency Thresholding

License: Protection

You can modify the rule latency threshold, the suspension time for suspended rules, and the number of consecutive threshold violations that must occur before suspending rules.

Rule latency thresholding suspends rules for the time specified by **Suspension Time** when the time rules take to process a packet exceeds **Threshold** for the consecutive number of times specified by **Consecutive Threshold Violations Before Suspending Rule**.

You can enable rule 134:1 to generate an event when rules are suspended, and rule 134:2 to generate an event when suspended rules are enabled. See Setting Rule States, page 27-19 for more information.

The following table further describes the options you can set to configure rule latency thresholding.

Table 11-6 Rule Latency Thresholding Options

Option	Description
Threshold	Specifies the time in microseconds that rules should not exceed when examining a packet. See the Minimum Rule Latency Threshold Settings table for recommended minimum threshold settings.
Consecutive Threshold Violations Before Suspending Rule	Specifies the consecutive number of times rules can take longer than the time set for **Threshold** to inspect packets before rules are suspended.
Suspension Time	Specifies the number of seconds to suspend a group of rules.

Many factors affect measurements of system performance, such as CPU speed, data rate, packet size, and protocol type. For this reason, Cisco recommends that you use the threshold settings in the following table until your own calculations provide you with settings tailored to your network environment.

Table 11-7 Minimum Rule Latency Threshold Settings

For this data rate...	Set threshold microseconds to at least...
1 Gbps	500
100 Mbps	1250
5 Mbps	5000

Determine the following when calculating your settings:

- average packets per second
- average microseconds per packet

Multiply the average microseconds per packet for your network by a significant safety factor to ensure that you do not unnecessarily suspend rules.

To configure rule latency thresholding:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy.**

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the Advanced tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Latency-Based Performance Settings**, then select the Rule Handling tab in the pop-up window that appears.

Step 5 You can configure any of the options in the Rule Latency Thresholding Options table.

See the Minimum Rule Latency Threshold Settings table for recommended minimum **Threshold** settings.

Step 6 Click **OK**.

You must save and apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Configuring Intrusion Performance Statistic Logging

License: Protection

You can configure the basic parameters of how devices monitor and report their own performance. This allows you to specify the intervals at which the system updates performance statistics on your devices by configuring the following options.

Sample time (seconds) and Minimum number of packets

When the number of seconds specified elapses between performance statistics updates, the system verifies it has analyzed the specified number of packets. If it has, the system updates performance statistics. Otherwise, the system waits until it analyzes the specified number of packets.

Troubleshooting Options: Log Session/Protocol Distribution

Support might ask you during a troubleshooting call to log protocol distribution, packet length, and port statistics.

Caution Changing the setting for this troubleshooting option will affect performance and should be done only with Support guidance.

Troubleshooting Options: Summary

Support might ask you during a troubleshooting call to configure the system to calculate the performance statistics only when the Snort® process is shut down or restarted. To enable this option, you must also enable the **Log Session/Protocol Distribution** troubleshooting option.

Caution Changing the setting for this troubleshooting option will affect performance and should be done only with Support guidance.

To configure basic performance statistics parameters:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy.**

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the Advanced tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Performance Settings**, then select the Performance Statistics tab in the pop-up window that appears.

Step 5 Modify the **Sample time** or **Minimum number of packets** as described above.

Step 6 Optionally, expand the **Troubleshoot Options** section and modify those options only if asked to do so by Support.

Step 7 Click **OK**.

You must save and apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Tuning File and Malware Inspection Performance and Storage

License: Protection or Malware

If you use file policies to perform file control or malware detection or blocking, you can set the options listed in the following table. Keep in mind that increasing the file sizes can affect the performance of the system.

Table 11-8 Advanced Access Control File and Malware Detection Options

Field	Description	Default Value	Range	Notes
Limit the number of bytes inspected when doing file type detection	Specify the number of bytes inspected when performing file type detection.	`1460` bytes, or the maximum segment size of a TCP packet	`0` - `4294967295` (4GB)	Set to `0` to remove the restriction. In most cases, the system can identify common file types using the first packet.
Do not calculate SHA-256 hash values for files larger than (in bytes)	Prevent the system from storing files larger than a certain size, performing a Collective Security Intelligence Cloud lookup on the files, or blocking the files if added to the custom detection list.	`10485760` (10MB)	`0` - `4294967295` (4GB)	Set to `0` to remove the restriction.
Allow file if cloud lookup for Block Malware takes longer than (seconds)	Specify how long the system will hold the last byte of a file that matches a **Block Malware** rule and that does not have a cached disposition, while malware cloud lookup occurs. If the time elapses without the system obtaining a disposition, the file passes. Dispositions of Unavailable are not cached.	`2` seconds	`0` - `30` seconds	Although this option accepts values of up to `30` seconds, Cisco recommends that you use the default value to avoid blocking traffic because of connection failures. Do **not** set this option to `0` without contacting Support.

To configure file and malware inspection performance and storage:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy.**

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the Advanced tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Files and Malware Settings**.

The Files and Malware Settings pop-up window appears.

Step 5 You can set any of the options in the Advanced Access Control File and Malware Detection Options table.

Step 6 Click **OK**.

You must save and apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.



Intelligent Application Bypass (IAB)

The following topics describe how to configure access control polices to use Intelligent Application Bypass:

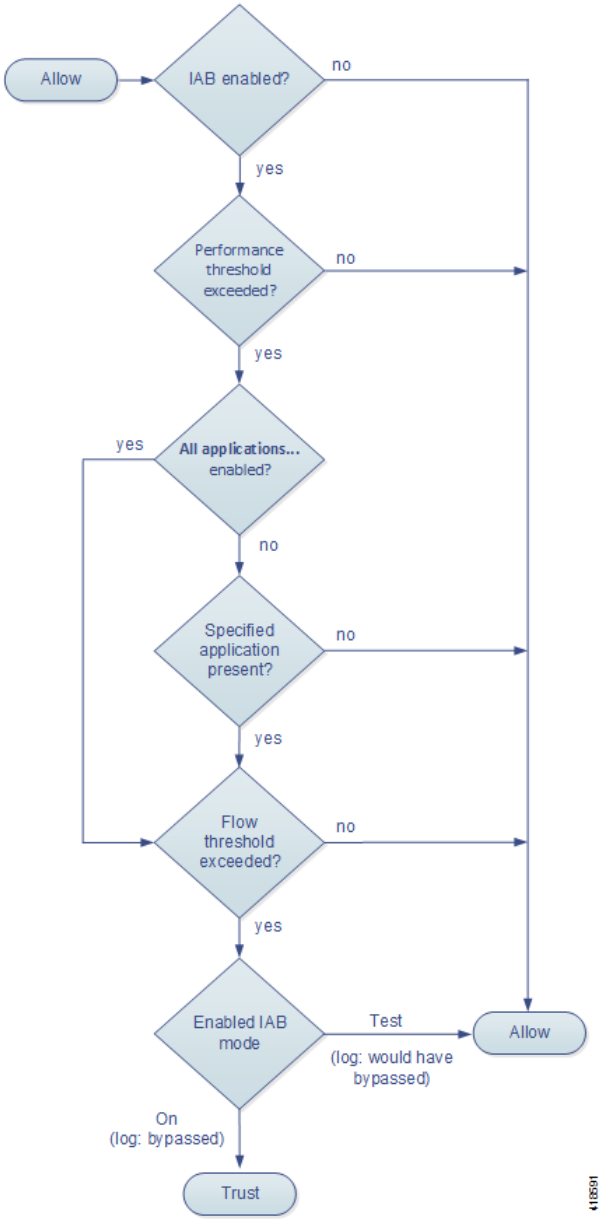
- Introduction to IAB, page 12-1
- IAB Options, page 12-2
- Configuring IAB, page 12-4
- IAB Logging and Analysis, page 12-5

Introduction to IAB

Intelligent Application Bypass (IAB) identifies applications that you trust to traverse your network without further inspection if performance and flow thresholds are exceeded. For example, if a nightly backup significantly impacts system performance, you can configure thresholds that, if exceeded, trust traffic generated by your backup application. Optionally, you can configure IAB so that, when an inspection performance threshold is exceeded, IAB trusts all traffic that exceeds any flow bypass threshold, regardless of the application type.

The system implements IAB on traffic allowed by access control rules or the access control policy's default action, before the traffic is subject to deep inspection. A test mode allows you to determine whether thresholds are exceeded and, if so, to identify the application flows that would have been bypassed if you had actually enabled IAB (called *bypass mode*).

The following graphic illustrates the IAB decision-making process:



IAB Options

State

Enables or disables IAB.

Performance Sample Interval

Specifies the time in seconds between IAB performance sampling scans, during which the system collects system performance metrics for comparison to IAB performance thresholds. A value of 0 disables IAB.

Bypassable Applications and Filters

This feature provides two mutually exclusive options:

Applications/Filters

Provides an editor where you can specify bypassable applications and sets of applications (filters) in essentially the same ways you specify application conditions in access control rules. See Controlling Application Traffic, page 8-2 for more information.

All applications including unidentified application

When an inspection performance threshold is exceeded, trusts all traffic that exceeds any flow bypass threshold, regardless of the application type.

Inspection Performance Thresholds

Inspection performance thresholds provide intrusion inspection performance limits that, if exceeded, trigger inspection of flow thresholds. IAB does not use inspection performance thresholds set to 0.



Note Inspection performance and flow bypass thresholds are disabled by default. You must enable at least one of each, and one of each must be exceeded for IAB to trust traffic. If you enable more than one inspection performance or flow bypass threshold, only one of each must be exceeded for IAB to trust traffic.

Drop Percentage

Average packets dropped as a percentage of total packets, when packets are dropped because of performance overloads caused by expensive intrusion rules, file policies, decompression, and so on. This does not refer to packets dropped by normal configurations such as intrusion rules. Note that specifying an integer greater than 1 activates IAB when the specified percentage of packets is dropped. When you specify 1, any percentage from 0 through 1 activates IAB. This allows a small number of packets to activate IAB.

Processor Utilization Percentage

Average percentage of processor resources used.

Package Latency

Average packet latency in microseconds.

Flow Rate

The rate at which the system processes flows, measured as the number of flows per second. Note that this option configures IAB to measure flow *rate*, not flow *count*.

Flow Bypass Thresholds

Flow bypass thresholds provide flow limits that, if exceeded, trigger IAB to trust bypassable application traffic in bypass mode or allow application traffic subject to further inspection in test mode. IAB does not use flow bypass thresholds set to 0.



Note Inspection performance and flow bypass thresholds are disabled by default. You must enable at least one of each, and one of each must be exceeded for IAB to trust traffic. If you enable more than one inspection performance or flow bypass threshold, only one of each must be exceeded for IAB to trust traffic.

Bytes per Flow

The maximum number of kilobytes a flow can include.

Packets per Flow

The maximum number of packets a flow can include.

Flow Duration

The maximum number of seconds a flow can remain open.

Flow Velocity

The maximum transfer rate in kilobytes per second.

Configuring IAB

Caution Not all deployments require IAB, and those that do might use it in a limited fashion. Do not enable IAB unless you have expert knowledge of your network traffic, especially application traffic, and system performance, including the causes of predictable performance issues. Before you run IAB in bypass mode, make sure that trusting the specified traffic does not expose you to risk.

To identify applications that you trust to traverse your network when thresholds are exceeded:

Step 1 In the access control policy editor, click the **Advanced** tab, then click the edit icon ( ) next to **Intelligent Application Bypass Settings**.

If a view icon ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck Inherit from base policy to enable editing.

Step 2 Configure IAB options:

- State—Turn IAB **Off** or **On**, or enable IAB in **Test** mode.
- Performance Sample Interval—Enter the time in seconds between IAB performance-sampling scans. If you enable IAB, even in test mode, enter a non-zero value. Entering 0 disables IAB.
- Bypassable Applications and Filters—Choose from:
  - Click the number of bypassed applications and filters and specify the applications whose traffic you want to bypass; specify in essentially the same ways you specify application conditions in access control rules. See Controlling Application Traffic, page 8-2 for more information.
  - Click **All applications including unidentified applications** so that, when an inspection performance threshold is exceeded, IAB trusts all traffic that exceeds any flow bypass threshold, regardless of the application type.
- Inspection Performance Thresholds—Click **Configure** and enter at least one threshold value.
- Flow Bypass Thresholds—Click **Configure** and enter at least one threshold value.

You must specify at least one inspection performance threshold and one flow bypass threshold; both must be exceeded for IAB to trust traffic. If you enter more than one threshold of each type, only one of each type must be exceeded. For detailed information, see IAB Options, page 12-2.

Step 3 Click **OK** to save IAB settings.

Step 4 Click **Save** to save the policy.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

IAB Logging and Analysis

IAB forces an end-of-connection event that logs bypassed flows and flows that would have been bypassed, regardless of whether you have enabled connection logging. Connection events indicate flows that are bypassed in bypass mode or that would have been bypassed in test mode. Custom dashboard widgets and reports based on connection events can display long-term statistics for bypassed and would-have-bypassed flows.

IAB Connection Events

Action

When **Reason** includes `Intelligent App Bypass`:

Allow—indicates that the applied IAB configuration was in test mode and traffic for the application specified by **Application Protocol** remains available for inspection.

Trust - indicates that the applied IAB configuration was in bypass mode and traffic for the application specified by **Application Protocol** has been trusted to traverse the network without further inspection.

Reason

`Intelligent App Bypass` indicates that IAB triggered the event in bypass or test mode.

Application Protocol

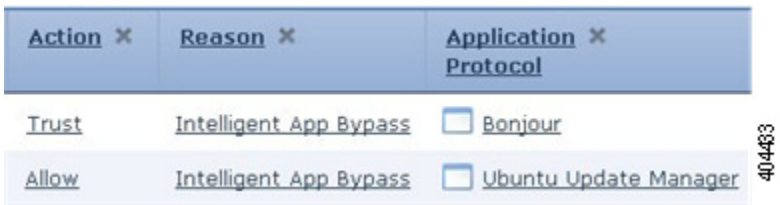
This field displays the application protocol that triggered the event.

Example

In the following truncated graphic, some fields are omitted. The graphic shows the **Action**, **Reason**, and **Application Protocol** fields for two connection events resulting from different IAB settings in two separate access control policies.

For the first event, the `Trust` action indicates that IAB was enabled in bypass mode and Bonjour protocol traffic was trusted to pass without further inspection.

For the second event, the `Allow` action indicates that IAB was enabled in test mode, so Ubuntu Update Manager traffic was subject to further inspection but would have been bypassed if IAB had been in bypass mode.



Example

In the following truncated graphic, some fields are omitted. The flow in the second event was both bypassed (**Action**: `Trust`; **Reason**: `Intelligent App Bypass`) and inspected by an intrusion rule (**Reason**: `Intrusion Monitor`). The `Intrusion Monitor` reason indicates that an intrusion rule set

to **Generate Events** detected but did not block an exploit during the connection. In the example, this happened before the application was detected. After the application was detected, IAB recognized the application as bypassable and trusted the flow.



IAB Custom Dashboard Widgets

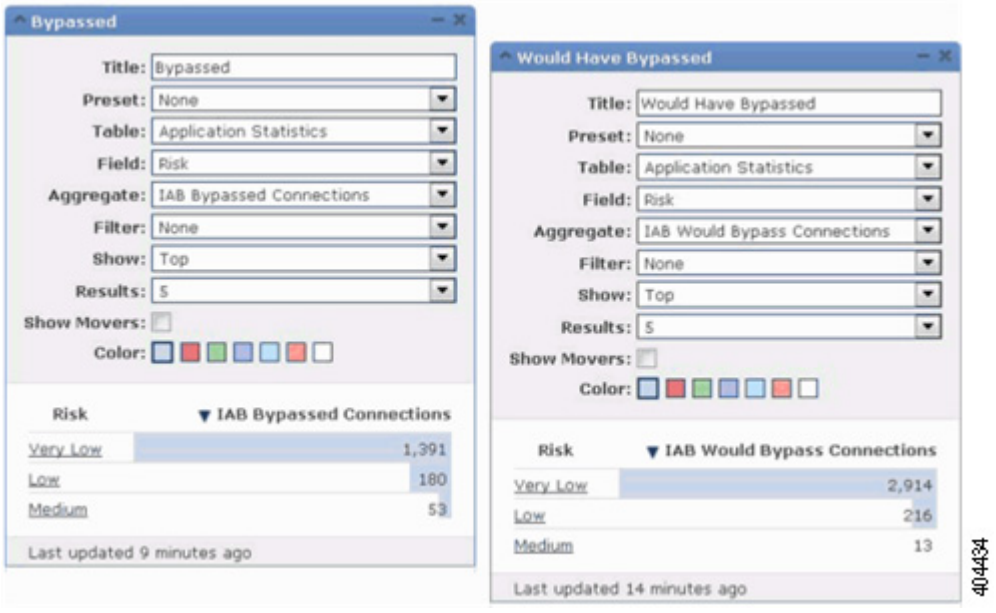
You can create a Custom Analysis dashboard widget to display long-term IAB statistics based on connection events. Specify the following when creating the widget:

- **Preset**: `None`
- **Table**: `Application Statistics`
- **Field**: `any`
- **Aggregate**: either of:
  - `IAB Bypassed Connections`
  - `IAB Would Bypass Connections`
- Filter: `any`

Examples

In the following Custom Analysis dashboard widget examples:

- The *Bypassed* example shows statistics for application traffic bypassed because the applications were specified as bypassable and IAB was enabled in bypass mode in the deployed access control policy.
- The *Would Have Bypassed* example shows statistics for application traffic that would have been bypassed because the applications were specified as bypassable and IAB was enabled in test mode in the deployed access control policy.



IAB Custom Reports

You can create a custom report to display long-term IAB statistics based on connection events. Specify the following when creating the report:

- **Table**: `Application Statistics`
- **Preset**: `None`
- **Filter**: `any`
- **X-Axis**: `any`
- **Y-Axis**: either of:
  - `IAB Bypassed Connections`
  - `IAB Would Bypass Connections`

Examples

The following graphic shows two abbreviated report examples:

- The *Bypassed* example shows statistics for application traffic bypassed because the applications were specified as bypassable and IAB was enabled in bypass mode in the deployed access control policy. The *Would Have Bypassed* example shows statistics for application traffic that would have been bypassed because the applications were specified as bypassable and IAB was enabled in test mode in the deployed access control policy.





Access Control Using Content Restriction

Major search engines and content delivery services provide features that allow you to restrict search results and website content. For example, schools use content restriction features to comply with the Children's Internet Protection Act (CIPA).

When implemented by search engines and content delivery services, you can enforce content restriction features only for individual browsers or users. The Firepower System allows you to extend these features to your entire network.

The system allows you to enforce:

- *Safe Search*—Supported in many major search engines, this service filters out explicit and adult-oriented content that particular environments (business, government, education, etc.) classify as objectionable. The system does not restrict a user's ability to access the home pages for supported search engines. Note that YouTube Restricted Mode is a subfeature of Safe Search.
- *YouTube EDU*—This service filters YouTube content for an educational environment. It allows schools to set access for educational content while limiting access to noneducational content. YouTube EDU is a different feature than YouTube Restricted Mode, which enforces restrictions on YouTube searches as part of Google's Safe Search feature. With YouTube EDU, users access the YouTube EDU home page, rather than the standard YouTube home page.

Content restriction features communicate the restricted status of a search or content query via an element in the request URI, an associated cookie, or a custom HTTP header element. You can configure access control rules to modify these elements as the system processes traffic.

Note that, to enforce content restriction, you must also enable an SSL policy, which impacts performance.

If you enable logging of connection events for these access control rules, the system logs related events with a **Reason** of `Content Restriction`.

The following topics describe how to enforce content restriction using access control rules:

- Using Access Control Rules to Enforce Content Restriction, page 13-1
- Safe Search Options for Access Control Rules, page 13-3
- YouTube EDU Options for Access Control Rules, page 13-3
- Content Restriction Rule Order, page 13-4

Using Access Control Rules to Enforce Content Restriction

License: Any

Caution To avoid rule preemption, position rules governing YouTube EDU above rules governing Safe Search in both SSL and access control policies. For more information, see Content Restriction Rule Order, page 13-4.

To enforce content restriction using access control rules:

Step 1 Create an SSL policy; see Creating a Basic SSL Policy, page 15-2.

Step 2 Add SSL rules for handling Safe Search and YouTube EDU traffic:

- Choose **Decrypt - Resign** as the **Action** for the rules. The system does not allow any other action for content restriction handling.
- In the **Applications** tab, add selections to the **Selected Applications and Filters** list:
  - Safe Search—Add the `safesearch supported` filter.
  - YouTube EDU—Search for "YouTube" in the **Available Applications** list, and add the resulting applications.

  For more information, see Controlling Encrypted Traffic Based on Application, page 17-8.

Step 3 Set rule positions for the SSL rules you added. Click and drag, or use the right-click menu to cut and paste.

Step 4 Create or edit an access control policy, and associate the SSL policy with the access control policy; see Associating Other Policies with Access Control, page 4-9.

Step 5 In the access control policy, add rules for handling Safe Search and YouTube EDU traffic, placing the Safe Search rule after the YouTube EDU rule:

- Choose **Allow** as the **Action** for the rules. The system does not allow any other action for content restriction handling.
- In the **Applications** tab, click the dimmed icon for either Safe Search ( ) or YouTube EDU ( ) , and set related options. These icons are disabled, rather than dimmed, if you choose any **Action** other than **Allow** for the rule.

Note You cannot enable Safe Search and YouTube EDU restrictions for the same access control rule.

- In the **Applications** tab, refine application selections in the **Selected Applications and Filters** list.

  In most cases, enabling Safe Search or YouTube EDU populates the **Selected Applications and Filters** list with the appropriate values. The system does not automatically populate the list if a Safe Search or YouTube application is already present in the list when you enable the feature. If applications do not populate as expected, manually add them as follows:
  - Safe Search—Add the `search engines` filter.
  - YouTube EDU—Search for "YouTube" in the **Available Applications** list, and add the resulting applications.

  For more information, see Adding an Application Condition to an Access Control Rule, page 8-5.

Step 6 Set rule positions for the access control rules you added. Click and drag, or use the right-click menu to cut and paste.

Step 7 Configure the **Block Response Page** that the system displays when it blocks restricted content; see Displaying a Custom Web Page for Blocked URLs, page 8-14.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Safe Search Options for Access Control Rules

The Firepower System supports Safe Search filtering for specific search engines only. For a list of supported search engines, see applications tagged `safesearch supported` in the **Applications** tab of the access control rule editor. For a list of unsupported search engines, see applications tagged `safesearch unsupported`.

When enabling Safe Search for an access control rule, set the following parameters:

Enable Safe Search

Enables Safe Search filtering for traffic that matches this rule.

Unsupported Search Traffic

Specifies the action you want the system to take when it processes traffic from unsupported search engines. If you choose **Block** or **Block with Reset**, you must also configure the HTTP response page that the system displays when it blocks restricted content; see Displaying a Custom Web Page for Blocked URLs, page 8-14.

YouTube EDU Options for Access Control Rules

When enabling YouTube EDU for an access control rule, set the following parameters:

Enable YouTube EDU

Enables YouTube EDU filtering for traffic that matches this rule.

Custom ID

Specifies the value that uniquely identifies a school or district network in the YouTube EDU initiative. YouTube provides this ID when a school or district registers for a YouTube EDU account.

Note If you check **Enable YouTube EDU**, you must enter a **Custom ID**. This ID is defined externally by YouTube. The system does not validate what you enter against the YouTube system. If you enter an invalid ID, YouTube EDU restrictions may not perform as expected.

Content Restriction Rule Order

To avoid rule preemption in both SSL and access control policies, position rules governing YouTube restriction above rules governing Safe Search restriction.

When you enable Safe Search for an access control rule, the system adds the search engine category to the **Selected Applications and Filters** list. This application category includes YouTube. As a result, YouTube traffic matches to the Safe Search rule unless YouTube EDU is enabled in a rule with a higher evaluation priority.

A similar rule preemption occurs if you position an SSL rule with the `safesearch supported` filter higher in the evaluation order than an SSL rule with specific YouTube application conditions.

For more information, see Ordering Rules to Improve Performance and Avoid Preemption, page 4-14.



Understanding Traffic Decryption

By default, the ASA FirePOWER module cannot inspect traffic encrypted with the Secure Socket Layer (SSL) protocol or its successor, the Transport Layer Security (TLS) protocol. As part of access control, the *SSL inspection* feature allows you to either block encrypted traffic without inspecting it, or inspect encrypted or decrypted traffic with access control. As the module handles encrypted sessions, it logs details about the traffic. The combination of inspecting encrypted traffic and analyzing encrypted session data allows greater awareness and control of the encrypted applications and traffic in your network.

SSL inspection is a policy-based feature. In the Firepower System, an access control policy is a master configuration that invokes subpolicies and other configurations, including an SSL policy. If you associate an SSL policy with access control, the system uses that SSL policy to handle encrypted sessions before it evaluates them with access control rules. If you do not configure SSL inspection, or your devices do not support it, access control rules handle all encrypted traffic.

Note that access control rules also handle encrypted traffic when your SSL inspection configuration allows it to pass. However, some access control rule conditions require unencrypted traffic, so encrypted traffic may match fewer rules. Also, by default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured. For more information, see Creating and Editing Access Control Rules, page 6-2 and Using the SSL Preprocessor, page 22-71.

If the module detects an SSL or TLS handshake over a TCP connection, it determines whether it can decrypt the detected traffic. If it cannot, it applies a configured action:

- block the encrypted traffic, and optionally reset the TCP connection
- not decrypt the encrypted traffic

If the module can decrypt the traffic, it blocks the traffic without further inspection, evaluates undecrypted traffic with access control, or decrypts it using one of the following methods:

- Decrypt with a known private key. When an external host initiates an SSL handshake with a server on your network, the system matches the exchanged server certificate with a server certificate previously uploaded to the appliance. It then uses the uploaded private key to decrypt the traffic.
- Decrypt by re-signing the server certificate. When a host on your network initiates an SSL handshake with an external server, the system re-signs the exchanged server certificate with a previously uploaded certificate authority (CA) certificate. It then uses the uploaded private key to decrypt the traffic.

Decrypted traffic is subject to the same traffic handling and analysis as originally unencrypted traffic: network, reputation, and user-based access control; intrusion detection and prevention; and advanced malware protection. If the system does not block the decrypted traffic post-analysis, it reencrypts the traffic before passing it to the destination host.

Note Certain SSL inspection actions, such as blocking traffic and decrypting outgoing traffic, modify the flow of traffic. ASA FirePOWER modules deployed inline can perform these actions. ASA FirePOWER modules deployed passively cannot affect the flow of traffic. However, these devices can still decrypt incoming traffic; see Example: Decrypting Traffic in a Passive Deployment, page 14-8 for more information.

For more information, see the following sections:

- SSL Handshake Processing, page 14-2
- SSL Inspection Requirements, page 14-5
- Analyzing SSL Inspection Appliance Deployments, page 14-7

SSL Handshake Processing

In this documentation, the term *SSL handshake* represents the two-way handshake that initiates encrypted sessions in both the SSL protocol and its successor protocol, TLS.

In a passive deployment, the Firepower System observes a copy of the handshake, but does not process the actual handshake. In an inline deployment, the Firepower System processes the SSL handshake, potentially modifying the ClientHello message and acting as a TCP proxy server for the session.

After the client establishes a TCP connection with the server (after it successfully completes the TCP three-way handshake), the managed device monitors the TCP session for any attempt to initiate an encrypted session. The SSL handshake establishes an encrypted session via the exchange of specialized packets between client and server. In the SSL and TLS protocols, these specialized packets are called *handshake messages*. The handshake messages communicate which encryption attributes both the client and server support:

- ClientHello—The client specifies multiple supported values for each encryption attribute.
- ServerHello—The server specifies a single supported value for each encryption attribute, which determines which encryption method the system uses during the secure session.

Although the data transmitted in the session is encrypted, the handshake messages are not.

After an SSL handshake completes, the managed device caches encrypted session data, which allows session resumption without requiring the full handshake. The managed device also caches server certificate data, which allows faster handshake processing in subsequent sessions.

ClientHello Message Handling

The client sends the ClientHello message to the server that acts as the packet destination if a secure connection can be established. The client sends the message to initiate the SSL handshake or in response to a Hello Request message from the destination server.

If you configure SSL inspection, when a managed device receives a ClientHello message, the system attempts to match the message to SSL rules that have the **Decrypt - Resign** action. The match relies on data from the ClientHello message and from cached server certificate data. Possible data includes:

Table 14-1 Data Availability for SSL Rule Conditions

SSL Rule Condition	Data Present In
Zones	ClientHello
Networks	ClientHello
VLAN Tags	ClientHello
Ports	ClientHello
Users	ClientHello
Applications	ClientHello (Server Name Indicator extension)
Categories	ClientHello (Server Name Indicator extension)
Certificate	server Certificate (potentially cached)
Distinguished Names	server Certificate (potentially cached)
Certificate Status	server Certificate (potentially cached)
Cipher Suites	ServerHello
Versions	ServerHello

If the ClientHello message does not match a Decrypt - Resign rule, the system does not modify the message. It then determines whether the message passes access control evaluation (which can include deep inspection). If the message passes, the system transmits it to the destination server.

If the message matches a Decrypt - Resign rule, the system modifies the ClientHello message as follows:

- Compression methods—Strips the `compression_methods` element, which specifies the compression methods the client supports. The Firepower System cannot decrypt compressed sessions. This modification reduces the Compressed Session type of undecryptable traffic.
- Cipher suites—Strips cipher suites from the `cipher_suites` element if the Firepower System does not support them. If the Firepower System does not support any of the specified cipher suites, the system transmits the original, unmodified element. This modification reduces the Unknown Cipher Suite and Unsupported Cipher Suite types of undecryptable traffic.
- Session identifiers—Strips any value from the `Session Identifier` element and the SessionTicket extension that does not match cached session data. If a ClientHello value matches cached data, an interrupted session can resume without the client and server performing the full SSL handshake. This modification increases the chances of session resumption and reduces the Session Not Cached type of undecryptable traffic.
- Elliptic curves—Strips elliptic curves from the Supported Elliptic Curves extension if the Firepower System does not support them. If the Firepower System does not support any of the specified elliptic curves, the managed device removes the extension and strips any related cipher suites from the `cipher_suites` element.
- ALPN extensions—Strips any value from the Application-Layer Protocol Negotiation (ALPN) extension that is unsupported in the Firepower System (for example, the SPDY and HTTP/2 protocols). This modification only occurs if the message matches an SSL rule associated with content restriction features. For more information, see Access Control Using Content Restriction, page 13-1.

- Other Extensions—Strips the Extended Master Secret, Next Protocol Negotiation (NPN), and TLS Channel IDs extensions.

Note The system performs these ClientHello modifications by default. If your SSL policy is configured correctly, this default behavior results in more frequent decryption of traffic. To tune the default behavior for your individual network, contact Support.

After the system modifies the ClientHello message, it determines whether the message passes access control evaluation (which can include deep inspection). If the message passes, the system transmits it to the destination server.

Direct communication between the client and server is no longer possible during the SSL handshake, because after message modification the Message Authentication Codes (MACs) computed by the client and server no longer match. For all subsequent handshake messages (and for the encrypted session once established), the managed device acts as a man-in-the-middle (MITM). It creates two SSL sessions, one between client and managed device, one between managed device and server. As a result, each session contains different cryptographic session details.

Note The cipher suites that the Firepower System can decrypt are frequently updated and do not correspond directly to the cipher suites you can use in SSL rule conditions. For the current list of decryptable cipher suites, contact Support.

ServerHello and Server Certificate Message Handling

The ServerHello message is the response to a ClientHello message in a successful SSL handshake.

After a managed device processes a ClientHello message and transmits it to the destination server, the server determines whether it supports the decryption attributes the client specified in the message. If it does not support those attributes, the server sends a handshake failure alert to the client. If it supports those attributes, the server sends the ServerHello message. If the agreed-upon key exchange method uses certificates for authentication, the server Certificate message immediately follows the ServerHello message.

When the managed device receives these messsages, it attempts to match them with SSL rules. These messages contain information that was absent from either the ClientHello message or the session data cache. Specifically, the system can potentially match these messages on Distinguished Names, Certificate Status, Cipher Suites, and Versions conditions.

If the messages do not match any SSL rules, the managed device performs the default action for the SSL policy. For more information, see Creating a Basic SSL Policy, page 15-2.

If the messages match an SSL rule, the managed device continues as appropriate:

Action: Monitor

The SSL handshake continues to completion. The managed device tracks and logs but does not decrypt encrypted traffic.

Action: Block or Block with Reset

The managed device blocks the SSL session. If appropriate, it also resets the TCP connection.

Action: Do Not Decrypt

The SSL handshake contintues to completion. The managed device does not decrypt the application data exchanged during the SSL session.

In rare cases, the system matches a ClientHello message to a Decrypt - Resign rule and modifies the message, but matches the related ServerHello message to a Do Not Decrypt rule. In those cases, the system resets the TCP connection to trigger a refreshed handshake from the client. The refreshed ClientHello message no longer matches the Decrypt - Resign rule, and the SSL session proceeds without decryption.

Action: Decrypt - Known Key

The managed device attempts to match the server Certificate data to a previously uploaded server certificate.

If it matches the certificate to a previously generated certificate, the SSL handshake continues to completion. The managed device uses the uploaded private key to decrypt and reencrypt the application data exchanged during the SSL session.

In rare cases, the system cannot match the server Certificate message to a previously generated certificate. For example, a server might change its certificate between the initial connection with the client and subsequent connections. In this case, the system blocks the SSL connection, so that the client reconnects and the system processes the handshake with the new certificate data.

Action: Decrypt - Resign

The managed device processes the server Certificate message and re-signs the exchanged server certificate with a previously uploaded certificate authority (CA) certificate. The SSL handshake continues to completion. The managed device then uses the uploaded private key to decrypt and reencrypt the application data exchanged during the SSL session.

While processing the ServerHello and Certificate messages, the managed device caches distinguished names and certificate data to allow faster handshake processing in both reestablished and subsequent SSL sessions.

SSL Inspection Requirements

License: feature dependent

How you deploy devices on your network, in addition to your configuration settings and licenses, influences the actions you can take to control and decrypt encrypted traffic.

SSL inspection requires public key certificates and paired private keys for certain features. You must upload certificates and paired private keys to the ASA FirePOWER module to decrypt and control traffic based on encryption session characteristics.

For more information, see the following sections:

- Deploying ASA FirePOWER Modules that Support SSL Inspection, page 14-6
- License Requirements for SSL Inspection, page 14-6
- Collecting Prerequisite Information to Configure SSL Rules, page 14-7

Deploying ASA FirePOWER Modules that Support SSL Inspection

License: Any

ASA FirePOWER modules configured and deployed inline can modify the flow of traffic. These devices can monitor, block, allow, and decrypt incoming and outgoing traffic.

ASA FirePOWER modules configured and deployed passively cannot affect the flow of traffic. They can only monitor, allow, and decrypt incoming traffic. Note that passive deployments do not support decrypting traffic encrypted with the ephemeral Diffie-Hellman (DHE) or the elliptic curve Diffie-Hellman (ECDHE) cipher suites.

Review your list of mapped actions, existing network deployment, and overall requirements to determine whether one or the other type of deployment better suits your organization. See Analyzing SSL Inspection Appliance Deployments, page 14-7 for more information.

License Requirements for SSL Inspection

License: feature dependent

Depending on your licenses, you can use a combination of criteria to determine how to handle encrypted traffic. The ASA FirePOWER module uses warning icons (⚠) and confirmation dialog boxes to designate unsupported features for your deployment. For details, hover your pointer over a warning icon.

You apply an SSL policy as part of an access control policy, and the access control policy inspects traffic decrypted by the SSL policy. See Access Control License and Role Requirements, page 4-2 for more information on access control licensing.

The following table explains the license requirements to apply an SSL policy as part of an access control policy.

Table 14-2 License Requirements for SSL Inspection

To apply an SSL policy that...	Licenses
handles encrypted traffic on the basis of zone, network, port, or SSL-related criteria	Any
handles encrypted traffic using geolocation data	Any
handles encrypted traffic using application or user criteria	Control
filters encrypted traffic using URL category and reputation data	URL Filtering

Collecting Prerequisite Information to Configure SSL Rules

License: feature-dependent

SSL inspection relies on a significant amount of supporting public key infrastructure (PKI) information. Consider your organization's traffic patterns to determine the matching rule conditions you can configure. Collect the information listed in the following table:

Table 14-3 SSL Rule Condition Prerequisites

To match on...	Collect the...
detected server certificates, including self-signed server certificates	server certificate
trusted server certificates	CA certificate
detected server certificate subject or issuer	server certificate subject DN or issuer DN

For more information, see Tuning Traffic Decryption Using SSL Rules, page 17-1.

Decide whether you want to not decrypt, block, monitor, or decrypt the encrypted traffic you match your rules against. Map these decisions to SSL rule actions, undecryptable traffic actions, and the SSL policy default action. If you want to decrypt traffic, see the following table:

Table 14-4 SSL Decryption Prerequisites

To decrypt...	Collect...
incoming traffic to a server you control	the server's certificate file and paired private key file
outgoing traffic to an external server	a CA certificate file and paired private key file You can also generate a CA certificate and private key.

For more information, see Using Rule Actions to Determine Encrypted Traffic Handling and Inspection, page 16-8.

After you have collected this information, upload it to the system and configure reusable objects. See Managing Reusable Objects, page 2-1 for more information.

Analyzing SSL Inspection Appliance Deployments

License: feature-dependent

This section presents several scenarios in which the Life Insurance Example, Inc. life insurance company (LifeIns) uses SSL inspection on encrypted traffic to help audit their processes. Based on their business processes, LifeIns plans to deploy:

- one ASA FirePOWER device in a passive deployment for the Customer Service department
- one ASA FirePOWER device in an inline deployment for the Underwriting Department

Customer Service Business Processes

LifeIns created a customer-facing website for their customers. LifeIns receives encrypted questions and requests regarding policies from prospective customers through their website and through e-mail. LifeIns's Customer Service department processes them and returns the requested information within 24 hours. Customer Service wants to expand its incoming contact metrics collection. LifeIns has an established internal audit review for Customer Service.

LifeIns also receives encrypted applications online. The Customer Service department processes the applications within 24 hours before sending the case file to the Underwriting department. Customer Service filters out any obvious false applications sent through the online form, which consumes a fair portion of their time.

Underwriting Business Processes

LifeIns's underwriters submit encrypted medical information requests online to the Medical Repository Example, LLC medical data repository (MedRepo). MedRepo reviews the requests and transmits the encrypted records to LifeIns within 72 hours. The underwriters subsequently underwrite an application and submit policy and rate decisions. Underwriting wants to expand its metrics collection.

Lately, an unknown source has been sending spoofed responses to LifeIns. Though LifeIns's underwriters receive training on proper Internet use, LifeIns's IT department first wants to analyze all encrypted traffic that takes the form of medical responses, then wants to block all spoof attempts.

LifeIns places junior underwriters on six-month training periods. Lately, these underwriters have been incorrectly submitting encrypted medical regulation requests to MedRepo's customer service department. MedRepo has submitted multiple complaints to LifeIns in response. LifeIns plans on extending their new underwriter training period to also audit underwriter requests to MedRepo.

For more information, see the following sections:

- Example: Decrypting Traffic in a Passive Deployment, page 14-8
- Example: Decrypting Traffic in an Inline Deployment, page 14-11

Example: Decrypting Traffic in a Passive Deployment

License: feature-dependent

LifeIns's business requirements state that Customer Service must:

- process all requests and applications within 24 hours
- improve its incoming contact metrics collection process
- identify and discard incoming false applications

Customer Service does not require additional audit review.

LifeIns plans to passively deploy a Customer Service device.

Traffic from an external network goes to LifeIns's router. The router routes traffic to the Customer Service department, and mirrors a copy of the traffic to the ASA FirePOWER module for inspection.

On the ASA FirePOWER module, a user in the Access Control and SSL Editor custom role configures SSL inspection to:

- log all encrypted traffic sent to the Customer Service department
- decrypt encrypted traffic sent using the online application form to Customer Service
- not decrypt all other encrypted traffic sent to Customer service, including traffic sent using the online request form

The user also configures access control to inspect the decrypted application form traffic for fake application data and log when fake data is detected.

In the following scenarios, the user submits an online form to Customer Service. The user's browser establishes a TCP connection with the server, then initiates an SSL handshake. The ASA FirePOWER module receives a copy of this traffic. The client and server complete the SSL handshake, establishing the encrypted session. Based on handshake and connection details, the system logs the connection and acts upon the copy of the encrypted traffic.

For more information, see the following:

- Monitoring Encrypted Traffic in a Passive Deployment, page 14-9
- Not Decrypting Encrypted Traffic in a Passive Deployment, page 14-9
- Inspecting Encrypted Traffic with a Private Key in a Passive Deployment, page 14-10

Monitoring Encrypted Traffic in a Passive Deployment

License: Any

For all SSL-encrypted traffic sent to Customer Service, the system logs the connection.

The following steps occur:

1. The user submits the plain text request (`info`). The client encrypts this (`AaBb`) and sends the encrypted traffic to Customer Service.
2. LifeIns's router receives the encrypted traffic and routes it to the Customer Service department server. It also mirrors a copy to the ASA FirePOWER module.
3. The Customer Service department server receives the encrypted information request (`AaBb`) and decrypts it to plain text (`info`).
4. The module does not decrypt the traffic.

   The access control policy continues to process the encrypted traffic and allows it. The module generates a connection event after the session ends.
5. The ASA FirePOWER module receives the connection event.

Not Decrypting Encrypted Traffic in a Passive Deployment

License: Any

For all SSL-encrypted traffic that contains requests about policies, the system allows the traffic without decrypting it and logs the connection.

The following steps occur:

1. The user submits the plain text request (`info`). The client encrypts this (`AaBb`) and sends the encrypted traffic to Customer Service.
2. LifeIns's router receives the encrypted traffic and routes it to the Customer Service department server. It also mirrors a copy to the ASA FirePOWER module.
3. The Customer Service department server receives the encrypted information request (`AaBb`) and decrypts it to plain text (`info`).
4. The ASA FirePOWER module does not decrypt the traffic.

   The access control policy continues to process the encrypted traffic and allows it. The module generates a connection event after the session ends.

5. The ASA FirePOWER module receives the connection event.

Inspecting Encrypted Traffic with a Private Key in a Passive Deployment

License: Any

For all SSL-encrypted traffic that contains application form data, the system decrypts the traffic and logs the connection.

Note In a passive deployment, if traffic is encrypted with either the DHE or ECDHE cipher suite, you cannot decrypt it with a known private key.

For traffic with legitimate application form information, the system logs the connection.

The following steps occur:

1. The user submits the plain text request (`form`). The client encrypts this (`AaBb`) and sends the encrypted traffic to Customer Service.
2. LifeIns's router receives the encrypted traffic and routes it to the Customer Service department server. It also mirrors a copy to the ASA FirePOWER module.
3. The Customer Service department server receives the encrypted information request (`AaBb`) and decrypts it to plain text (`form`).
4. The ASA FirePOWER module uses the session key obtained with the uploaded known private key to decrypt the encrypted traffic to plain text (`form`).

   The access control policy continues to process the decrypted traffic and does not find fake application information. The module generates a connection event after the session ends.
5. The ASA FirePOWER module receives a connection event with information about the encrypted and decrypted traffic.

In contrast, if the decrypted traffic contains fake application data, the system logs the connection and the fake data.

The following steps occur:

1. The user submits the plain text request (`fake`). The client encrypts this (`CcDd`) and sends the encrypted traffic to Customer Service.
2. LifeIns's router receives the encrypted traffic and routes it to the Customer Service department server. It also mirrors a copy to the device.
3. The Customer Service department server receives the encrypted information request (`CcDd`) and decrypts it to plain text (`fake`).
4. The ASA FirePOWER module uses the session key obtained with the uploaded known private key to decrypt the encrypted traffic to plain text (`fake`).

   The access control policy continues to process the decrypted traffic and finds fake application information. The module generates an intrusion event. After the session ends, it generates a connection event.
5. The ASA FirePOWER module receives a connection event with information about the encrypted and decrypted traffic, and an intrusion event for the fake application data.

Example: Decrypting Traffic in an Inline Deployment

License: feature-dependent

LifeIns's business requirements state that Underwriting must:

- audit new and junior underwriters, verifying that their information requests to MedRepo comply with all applicable regulations
- improve its underwriting metrics collection process
- examine all requests that appear to come from MedRepo, then drop any spoofing attempts
- drop all improper regulatory requests to MedRepo's Customer Service department from the Underwriting department
- not audit senior underwriters

LifeIns plans to deploy a device in an inline deployment for the Underwriting department.

Traffic from MedRepo's network goes to MedRepo's router. It routes traffic to LifeIns's network. The device receives the traffic, passes allowed traffic to LifeIns's router, and sends events to the ASA FirePOWER module. LifeIns's router routes traffic to the destination host.

On the ASA FirePOWER module, a user configures SSL inspection to:

- log all encrypted traffic sent to the Underwriting department
- block all encrypted traffic incorrectly sent from LifeIns's underwriting department to MedRepo's customer service department
- decrypt all encrypted traffic sent from MedRepo to LifeIns's underwriting department, and from LifeIns's junior underwriters to MedRepo's requests department
- not decrypt encrypted traffic sent from the senior underwriters

The user also configures access control to inspect decrypted traffic with a custom intrusion policy and:

- block decrypted traffic if it contains a spoof attempt, and log the spoof attempt
- block decrypted traffic that contains information not compliant with regulations, and log the improper information
- allow all other encrypted and decrypted traffic

The system reencrypts allowed decrypted traffic before sending it to the destination host.

In the following scenarios, the user submits information online to a remote server. The user's browser establishes a TCP connection with the server, then initiates an SSL handshake. The module receives this traffic; based on handshake and connection details, the system logs the connection and acts on the traffic. If the system blocks the traffic, it also closes the TCP connection. Otherwise, the client and server complete the SSL handshake, establishing the encrypted session.

For more information, see the following:

- Monitoring Encrypted Traffic in an Inline Deployment, page 14-12
- Allowing Specific Users' Encrypted Traffic in an Inline Deployment, page 14-12
- Blocking Encrypted Traffic in an Inline Deployment, page 14-12
- Inspecting Encrypted Traffic with a Private Key in an Inline Deployment, page 14-13
- Inspecting Specific Users' Encrypted Traffic with a Re-signed Certificate in an Inline Deployment, page 14-14

Monitoring Encrypted Traffic in an Inline Deployment

License: Any

For all SSL-encrypted traffic sent to and from the Underwriting department, the system logs the connection.

The following steps occur:

1. The user submits the plain text request (`help`). The client encrypts this (`AaBb`) and sends the encrypted traffic to MedRepo's Requests department server.
2. LifeIns's router receives the encrypted traffic and routes it to the Requests department server.
3. The ASA FirePOWER module does not decrypt the traffic.

   The access control policy continues to process the encrypted traffic and allows it, then generates a connection event after the session ends.
4. The external router receives the traffic and routes it to the Requests department server.
5. The Underwriting department server receives the encrypted information request (`AaBb`) and decrypts it to plain text (`help`).
6. The ASA FirePOWER module receives the connection event.

Allowing Specific Users' Encrypted Traffic in an Inline Deployment

License: Control

For all SSL-encrypted traffic originating from the senior underwriters, the system allows the traffic without decrypting it and logs the connection.

The following steps occur:

1. The user submits the plain text request (`help`). The client encrypts this (`AaBb`) and sends the encrypted traffic to MedRepo's Requests department server.
2. LifeIns's router receives the encrypted traffic and routes it to the Requests department server.
3. The ASA FirePOWER module does not decrypt this traffic.

   The access control policy continues to process the encrypted traffic and allows it, then generates a connection event after the session ends.
4. The external router receives the traffic and routes it to the Requests department server.
5. The Requests department server receives the encrypted information request (`AaBb`) and decrypts it to plain text (`help`).
6. The ASA FirePOWER module receives the connection event.

Blocking Encrypted Traffic in an Inline Deployment

License: Any

For all SMTPS email traffic improperly sent from LifeIns's underwriting department to MedRepo's Customer Service department, the system blocks the traffic during the SSL handshake without further inspection and logs the connection.

The following steps occur:

1. Having received the request to establish an SSL handshake from a client's browser, the Customer Service department server sends the server certificate (`cert`) as the next step in the SSL handshake to the LifeIns underwriter.
2. MedRepo's router receives the certificate and routes it to the LifeIns underwriter.
3. The ASA FirePOWER module blocks the traffic without further inspection and ends the TCP connection. It generates a connection event.
4. The internal router does not receive the blocked traffic.
5. The underwriter does not receive the blocked traffic.
6. The ASA FirePOWER module receives the connection event.

Inspecting Encrypted Traffic with a Private Key in an Inline Deployment

License: Any

For all SSL-encrypted traffic sent from MedRepo to LifeIns's underwriting department, the system uses an uploaded server private key to obtain session keys, then decrypts the traffic and logs the connection. Legitimate traffic is allowed and reencrypted before being sent to the Underwriting department.

The following steps occur:

1. The user submits the plain text request (`stats`). The client encrypts this (`AaBbC`) and sends the encrypted traffic to the Underwriting department server.
2. The external router receives the traffic and routes it to the Underwriting department server.
3. The ASA FirePOWER module uses the session key obtained with the uploaded known private key to decrypt this traffic to plain text (`stats`).

   The access control policy continues to process the decrypted traffic with the custom intrusion policy and does not find a spoof attempt. The device passes the encrypted traffic (`AaBbC`), then generates a connection event after the session ends.
4. The internal router receives the traffic and routes it to the Underwriting department server.
5. The Underwriting department server receives the encrypted information (`AaBbC`) and decrypts it to plain text (`stats`).
6. The ASA FirePOWER module receives the connection event with information about the encrypted and decrypted traffic.

In contrast, any decrypted traffic that is a spoof attempt is dropped. The system logs the connection and the spoof attempt.

The following steps occur:

1. The user submits the plain text request (`spoof`), altering the traffic to appear to originate from MedRepo, LLC. The client encrypts this (`FfGgH`) and sends the encrypted traffic to the Underwriting department server.
2. The ASA FirePOWER module uses the session key obtained with the uploaded known private key to decrypt this traffic to plain text (`spoof`).

   The access control policy continues to process the decrypted traffic with the custom intrusion policy and finds a spoof attempt. The ASA FirePOWER module blocks the traffic, then generates an intrusion event. It generates a connection event after the session ends.
3. The internal router does not receive the blocked traffic.

4. The Underwriting department server does not receive the blocked traffic.
5. The ASA FirePOWER module receives a connection event with information about the encrypted and decrypted traffic, and an intrusion event for the spoofing attempt.

Inspecting Specific Users' Encrypted Traffic with a Re-signed Certificate in an Inline Deployment

License: Control

For all SSL-encrypted traffic sent from the new and junior underwriters to MedRepo's requests department, the system uses a re-signed server certificate to obtain session keys, then decrypts the traffic and logs the connection. Legitimate traffic is allowed and reencrypted before being sent to MedRepo.



Note When decrypting traffic in an inline deployment by re-signing the server certificate, the ASA FirePOWER module acts as a man-in-the-middle. It creates two SSL sessions, one between client and ASA FirePOWER module, one between ASA FirePOWER module and server. As a result, each session contains different cryptographic session details.

The following steps occur:

1. The user submits the plain text request (`help`). The client encrypts this (`AaBb`) and sends the encrypted traffic to the Requests department server.
2. The internal router receives the traffic and routes it to the Requests department server.
3. The ASA FirePOWER module uses the session key obtained with a re-signed server certificate and private key to decrypt this traffic to plain text (`help`).

   The access control policy continues to process the decrypted traffic with the custom intrusion policy and does not find an improper request. The module reencrypts the traffic (`CcDd`), allowing it to pass. It generates a connection event after the session ends.
4. The external router receives the traffic and routes it to the Requests department server.
5. The Requests department server receives the encrypted information (`CcDd`) and decrypts it to plain text (`help`).
6. The ASA FirePOWER module receives the connection event with information about the encrypted and decrypted traffic.



Note Traffic encrypted with a re-signed server certificate causes client browsers to warn that the certificate is not trusted. To avoid this, add the CA certificate to the organization's domain root trusted certificates store or the client trusted certificates store.

In contrast, any decrypted traffic that contains information that does not meet regulatory requirements is dropped. The system logs the connection and the non-conforming information.

The following steps occur:

1. The user submits the plain text request (`regs`), which does not comply with regulatory requirements. The client encrypts this (`EeFf`) and sends the encrypted traffic to the Requests department server.
2. The internal router receives the traffic and routes it to the Requests department server.

3. The ASA FirePOWER module uses the session key obtained with a re-signed server certificate and private key to decrypt this traffic to plain text (`regs`).

   The access control policy continues to process the decrypted traffic with the custom intrusion policy and finds an improper request. The module blocks the traffic, then generates an intrusion event. It generates a connection event after the session ends.

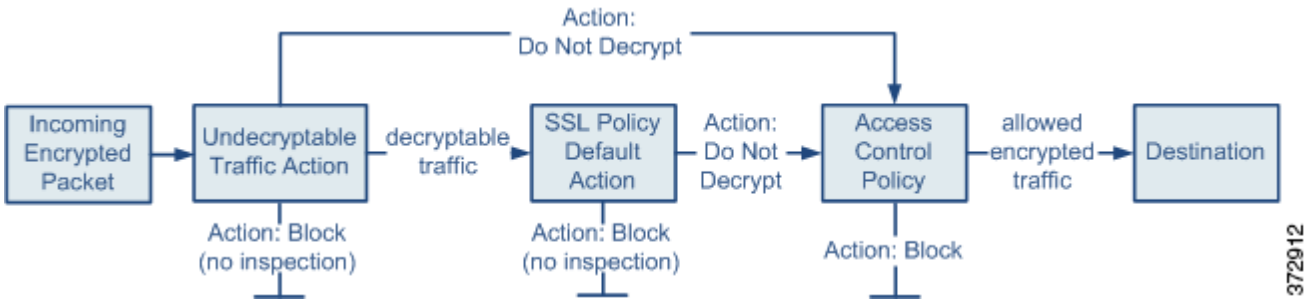
4. The external router does not receive the blocked traffic.
5. The Requests department server does not receive the blocked traffic.
6. The ASA FirePOWER module receives a connection event with information about the encrypted and decrypted traffic, and an intrusion event for the improper request.



Getting Started with SSL Policies

An *SSL policy* determines how the system handles encrypted traffic on your network. You can configure one or more SSL policies. You associate an SSL policy with an access control policy, then apply the access control policy. When the ASA FirePOWER module detects a TCP handshake, the access control policy first handles and inspects the traffic. If it subsequently identifies an SSL-encrypted session over the TCP connection, the SSL policy takes over, handling and decrypting the encrypted traffic. You can have one currently applied SSL policy.

The simplest SSL policy, as shown in the following diagram, directs the device where it is applied to handle encrypted traffic with a single default action. You can set the default action to block decryptable traffic without further inspection, or inspect undecrypted decryptable traffic with access control. The system can then either allow or block the encrypted traffic. If the ASA FirePOWER module detects undecryptable traffic, it either blocks the traffic without further inspection or does not decrypt it, inspecting it with access control.



This chapter explains how to create and apply a simple SSL policy. It also contains basic information on managing SSL policies: editing, updating, comparing, and so on. For more information, see:

- Creating a Basic SSL Policy, page 15-2
- Editing an SSL Policy, page 15-6
- Applying Decryption Settings Using Access Control, page 15-8
- Generating a Report of Current Traffic Decryption Settings, page 15-9
- Comparing SSL Policies, page 15-10

A more complex SSL policy can handle different types of undecryptable traffic with different actions, control traffic based on whether a certificate authority (CA) issued or trusts the encryption certificate, and use SSL rules to exert granular control over encrypted traffic logging and handling. These rules can be simple or complex, matching and inspecting encrypted traffic using multiple criteria. After you create a basic SSL policy, see the following chapters for more information on tailoring it to your deployment:

- Managing Reusable Objects, page 2-1 describes how to configure reusable public key infrastructure (PKI) objects and other SSL inspection-related objects to enhance encrypted traffic control and decrypt traffic.
- Logging Connections in Network Traffic, page 36-1 describes how to configure logging for encrypted traffic, whether decryptable or undecryptable.
- Applying Decryption Settings Using Access Control, page 15-8 describes how to associate an SSL policy with an access control policy.
- Getting Started with Access Control Policies, page 4-1 describes how to apply an access control policy to a device.
- Tuning Traffic Flow Using Access Control Rules, page 6-1 describes how to configure access control rules to inspect decrypted traffic.
- Getting Started with SSL Rules, page 16-1 describes how to configure SSL rules to handle and log encrypted traffic.
- Tuning Traffic Decryption Using SSL Rules, page 17-1 describes how to configure SSL rule conditions to better match specific encrypted traffic.

Creating a Basic SSL Policy

License: Any

When you create a new SSL policy you must, at minimum, give it a unique name and specify a policy default action. You have the following options when selecting a default action for a new policy:

- **Do not decrypt** creates a policy with the Do not decrypt default action.
- **Block** creates a policy with the Block default action.
- **Block with reset** creates a policy with the Block with reset default action.

After you create the SSL policy, you can modify the default action. For guidance on choosing a default action, see Setting Default Handling and Inspection for Encrypted Traffic, page 15-3.

The new SSL policy also contains default actions for traffic the system cannot decrypt: either it inherits the default action you just selected for undecryptable traffic, blocks it, or does not decrypt the traffic and inspects it with access control. You can modify the undecryptable traffic actions after you create the SSL policy. For guidance on selecting undecryptable traffic actions, see Setting Default Handling for Undecryptable Traffic, page 15-4

On the SSL policy page (**Configuration > ASA FirePOWER Configuration > Policies > SSL**) you can view all your current SSL policies by name with optional description. Options on this page allow you to compare policies, create a new policy, copy a policy, view a report that lists all of the most recently saved settings in each policy, edit a policy, or delete a policy.

The following table describes the actions you can take to manage your policies on the SSL Policy page:

Table 15-1 SSL Policy Management Actions

To...	You can...
create a new SSL policy	click **New Policy**. See Creating a Basic SSL Policy, page 15-2 for more information.
modify the settings in an existing SSL policy	click the edit icon ( ). See Editing an SSL Policy, page 15-6 for more information.

Table 15-1 SSL Policy Management Actions (continued)

To...	You can...
compare SSL policies	click **Compare Policies**. See Comparing SSL Policies, page 15-10 for more information.
copy an SSL policy	click the copy icon ( ). See Editing an SSL Policy, page 15-6 for more information on editing a copied policy.
view a PDF report that lists the current configuration settings in an SSL policy	click the report icon ( ). See Generating a Report of Current Traffic Decryption Settings, page 15-9 for more information.
delete an SSL policy	click the delete icon ( ), then click **OK**. When prompted whether to continue, you are also informed if another user has unsaved changes in the policy.

To create an SSL policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > SSL**.

The SSL Policy page appears.

Step 2 Give the policy a unique **Name** and, optionally, a **Description**.

You can use all printable characters, including spaces and special characters.

Step 3 Specify the **Default Action**.

Note that you can modify your selected default action after you create your SSL policy. See Setting Default Handling and Inspection for Encrypted Traffic, page 15-3 for more information.

Step 4 Click **Store ASA FirePOWER Changes**.

The SSL Policy Editor page appears. See Editing an SSL Policy, page 15-6 for more information.

Setting Default Handling and Inspection for Encrypted Traffic

License: Any

The default action for an SSL policy determines how the system handles decryptable encrypted traffic that does not match any non-Monitor rule in the policy. When you apply an SSL policy that does not contain any SSL rules, the default action determines how all decryptable traffic on your network is handled. See Setting Default Handling for Undecryptable Traffic, page 15-4 for more information on how the system handles undecryptable encrypted traffic.

The following table lists the default actions you can choose, as well as their effect on encrypted traffic. Note that the system does not perform any kind of inspection on encrypted traffic blocked by the default action.

Table 15-2 SSL Policy Default Actions

Default Action	Effect on Encrypted Traffic
Block	block the SSL session without further inspection
Block with reset	block the SSL session without further inspection and reset the TCP connection
Do not decrypt	inspect the encrypted traffic with access control

When you first create an SSL policy, logging connections that are handled by the default action is disabled by default. You can change this, as well as the default action itself, after you create the policy.

The following procedure explains how to set the default action for an SSL policy while editing the policy. See Editing an SSL Policy, page 15-6 for the complete procedure for editing an SSL policy.

To set the default action of an SSL policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > SSL**.

The SSL policy page appears.

Step 2 Click the edit icon ( ) next to the SSL policy you want to configure.

The SSL policy editor appears.

Step 3 Select a **Default Action**. See the SSL Policy Default Actions table for more information.

Step 4 Configure logging options for the default action as described in Logging Decryptable Connections with SSL Rules, page 36-14.

Step 5 Click **Store ASA FirePOWER Changes**.

The SSL Policy Editor page appears. See Editing an SSL Policy, page 15-6 for more information.

Setting Default Handling for Undecryptable Traffic

License: Any

You can set undecryptable traffic actions at the SSL policy level to handle certain types of encrypted traffic the system cannot decrypt or inspect. When you apply an SSL policy that does not contain any SSL rules, the undecryptable traffic actions determine how all undecryptable encrypted traffic on your network is handled.

Depending on the type of undecryptable traffic, you can choose to:

- block the connection
- block the connection, then reset it
- inspect the encrypted traffic with access control
- inherit the default action from the SSL policy

The following table describes the undecryptable traffic types:

Table 15-3 Undecryptable Traffic Types

Type	Description	Default Action	Available Actions
Compressed Session	The SSL session applies a data compression method.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
SSLv2 Session	The session is encrypted with SSL version 2. Note that traffic is decryptable if the client hello message is SSL 2.0, and the remainder of the transmitted traffic is SSL 3.0.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Unknown Cipher Suite	The system does not recognize the cipher suite.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Unsupported Cipher Suite	The system does not support decryption based on the detected cipher suite.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Session not cached	The SSL session has session reuse enabled, the client and server reestablished the session with the session identifier, and the system did not cache that session identifier.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Handshake Errors	An error occurred during SSL handshake negotiation.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Decryption Errors	An error occurred during traffic decryption.	Block	Block Block with Reset

When you first create an SSL policy, logging connections that are handled by the default action is disabled by default. Because the logging settings for the default action also apply to undecryptable traffic handling, logging connections handled by the undecryptable traffic actions is disabled by default. For more information on configuring default logging, see Logging Decryptable Connections with SSL Rules, page 36-14.

Note The system cannot decrypt traffic if an HTTP proxy is positioned between a client and your device, and the client and server establish a tunneled SSL connection using the CONNECT HTTP method. The **Handshake Errors** undecryptable action determines how the system handles this traffic. See Decrypt Actions: Decrypting Traffic for Further Inspection, page 16-10 for more information.

Note that if your browser uses certificate pinning to verify a server certificate, you cannot decrypt this traffic by re-signing the server certificate. Because you can still inspect this traffic with access control, it is not handled by the undecryptable traffic actions. If you want to allow this traffic, configure an SSL rule with the Do not decrypt action to match the server certificate common name or distinguished name.

To set the default handling for undecryptable traffic:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > SSL**.

The SSL Policy page appears.

Step 2 Click the edit icon ( ) next to the SSL policy you want to configure.

The SSL policy editor appears.

Step 3 Select the **Undecryptable Actions** tab.

The Undecryptable Actions tab appears.

Step 4 For each field, select the action you want to take on the type of undecryptable traffic, or if you want to apply the SSL policy's default action. See the SSL Policy Default Actions table for more information.

Step 5 Click **Store ASA FirePOWER Changes**.

You must apply the associated access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Editing an SSL Policy

License: Any

On the SSL policy editor, you can configure your policy and organize SSL rules. To configure an SSL policy, you must give the policy a unique name and specify a default action. You can also:

- add, edit, delete, enable, and disable SSL rules
- add trusted CA certificates
- determine the handling for encrypted traffic the system cannot decrypt
- log traffic that is handled by the default action and undecryptable traffic actions

After you create or modify an SSL policy, you can associate it with an access control policy, then apply the access control policy. You can also create custom user roles that allow you to assign different permissions to different users for configuring, organizing, and applying policies.

The following table summarizes the configuration actions you can take on the SSL policy editor.

Table 15-4 SSL Policy Configuration Actions

To...	You can...
modify the policy name or description	click the name or description field, delete any characters as needed, then type the new name or description.
set the default action	find more information at Setting Default Handling and Inspection for Encrypted Traffic, page 15-3.
set default handling for undecryptable traffic	find more information at Setting Default Handling for Undecryptable Traffic, page 15-4.

Table 15-4 SSL Policy Configuration Actions (continued)

To...	You can...
log connections for the default action and undecryptable traffic actions	find more information at Logging Decryptable Connections with SSL Rules, page 36-14.
add trusted CA certificates	find more information at Trusting External Certificate Authorities, page 17-21.
assign different rights to different users	find more information at Collecting Prerequisite Information to Configure SSL Rules, page 14-7.
save your policy changes	click **Save**.
cancel your policy changes	click **Cancel**, then, if you have made changes, click **OK**.
add a rule to a policy	click **Add Rule**. See Understanding and Creating SSL Rules, page 16-4 for more information. **Tip** You can also right-click a blank area in the row for a rule and select **Insert new rule**.
edit an existing rule	click the edit icon ( ) next to the rule. See Understanding and Creating SSL Rules, page 16-4 for more information. **Tip** You can also right-click the rule and select **Edit**.
delete a rule	click the delete icon ( ) next to the rule, then click **OK**. **Tip** You can also right-click a blank area in the row for a selected rule, select **Delete**, then click **OK** to delete one or more selected rules.
enable or disable an existing rule	right-click a selected rule, select **State**, then select **Disable** or **Enable**. Disabled rules are grayed and marked `(disabled)` beneath the rule name.
display the configuration page for a specific rule attribute	click the name, value, or icon in the column for the condition on the row for the rule. For example, click the name or value in the **Source Networks** column to display the Networks page for the selected rule. See Tuning Traffic Decryption Using SSL Rules, page 17-1 for more information.

When you change your configuration, a message indicates that you have unsaved changes. To retain your changes, you must save the policy before exiting the policy editor. If you attempt to exit the policy editor without saving your changes, you are cautioned that you have unsaved changes; you can then discard your changes and exit the policy, or return to the policy editor.

To protect the privacy of your session, after sixty minutes of inactivity on the policy editor, changes to your policy are discarded and you are returned to the SSL Policy page. After the first thirty minutes of inactivity, a message appears and updates periodically to provide the number of minutes remaining before changes are discarded. Any activity on the page cancels the timer.

When multiple users edit the same policy concurrently, a message on the policy editor identifies other users who have unsaved changes. Any user who attempts to save changes is cautioned that his changes will overwrite changes by other users. When the same policy is saved by multiple users, the last saved changes are retained.

To edit an SSL policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > SSL**.

The SSL Policy page appears.

Step 2 You have the following choices:

- To configure your policy, you can take any of the actions summarized in the SSL Policy Configuration Actions table.
- To organize rules in your policy, you can take any of the actions described in Managing SSL Rules in a Policy, page 16-12.

Step 3 Save or discard your configuration. You have the following choices:

- To save your changes and continue editing, click **Store ASA FirePOWER Changes**.
- To discard your changes, click **Cancel** and, if prompted, click **OK**.

  Your changes are discarded and the SSL Policy page appears.

Applying Decryption Settings Using Access Control

License: Any

After making any changes to an SSL policy, you must apply the access control policy it is associated with. For more information, see Deploying Configuration Changes, page 4-11.

Keep the following points in mind when applying SSL policies:

- You cannot delete an SSL policy that has been applied or is currently applying.
- Applying an access control policy automatically applies the associated SSL policy. You cannot apply an SSL policy independently.

Note In a passive deployment, the system cannot influence the flow of traffic. If you attempt to apply an access control policy that references an SSL policy that blocks encrypted traffic, or that is configured to decrypt traffic by re-signing the server certificate, the system displays a warning. Also, passive deployments do not support decrypting traffic encrypted with the ephemeral Diffie-Hellman (DHE) or the elliptic curve Diffie-Hellman (ECDHE) cipher suites.

To associate an SSL policy with an access control policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to configure.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

Advanced settings for the access control policy appear.

Step 4 Click the edit icon ( ) next to General Settings.

The General Settings pop-up window appears.

Step 5 Select an SSL policy from the **SSL Policy to use for inspecting encrypted connections** drop-down.

Step 6 Click **OK**.

Advanced settings for the access control policy appear.

Step 7 Click **Store ASA FirePOWER Changes**.

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Generating a Report of Current Traffic Decryption Settings

License: Any

An SSL policy report is a record of the policy and rules configuration at a specific point in time. You can use the report for auditing purposes or to inspect the current configuration.

Tip You can also generate an SSL comparison report that compares a policy with the currently applied policy or with another policy. For more information, see Comparing SSL Policies, page 15-10.

An SSL policy report contains the sections described in the following table.

Table 15-5 SSL Policy Report Sections

Section	Description
Title Page	Identifies the name of the policy report, the date and time the policy was last modified, and the name of the user who made that modification.
Table of Contents	Describes the contents of the report.
Policy Information	Provides the name and description of the policy, the name of the user who last modified the policy, and the date and time the policy was last modified.
Default Action	Provides the default action.
Default Logging	Provides the default connection logging settings.
Rules	Provides the rule action and conditions for each rule in the policy, by rule category.
Trusted CA Certificates	Provides the CA certificates that are automatically trusted if detected traffic is encrypted using these certificates or other certificates within the chain of trust.
Undecryptable Actions	Provides the action taken on detected types of traffic that cannot be decrypted.
Referenced Objects	Provides the name and configuration of all individual objects and group objects used in the policy, by type of condition (networks, ports, and so on) where the object is configured.

To view an SSL policy report:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > SSL.**

The SSL Policy page appears.

Step 2 Click the report icon ( ) next to the policy for which you want to generate a report. Remember to save any changes before you generate an SSL policy report; only saved changes appear in the report.

The system generates the report. Depending on your browser settings, the report may appear in a pop-up window, or you may be prompted to save the report to your computer.

Comparing SSL Policies

License: Any

To review policy changes for compliance with your organization's standards or to optimize system performance, you can examine the differences between two SSL policies. You can compare any two policies or the currently applied policy with another policy. Optionally, after you compare, you can then generate a PDF report to record the differences between the two policies.

There are two tools you can use to compare policies:

- The comparison view displays only the differences between two policies in a side-by-side format. The name of each policy appears in the title bar on the left and right sides of the comparison view except when you select **Running Configuration**, in which case a blank bar represents the currently active policy.

  You can use this to view and navigate both policies on the web interface, with their differences highlighted.

- The comparison report creates a record of only the differences between two policies in a format similar to the policy report, but in PDF format.

  You can use this to save, copy, print, and share your policy comparisons for further examination.

For more information on understanding and using the policy comparison tools, see:

- Using the SSL Policy Comparison View, page 15-10
- Using the SSL Policy Comparison Report, page 15-11

Using the SSL Policy Comparison View

License: Any

The comparison view displays both policies in a side-by-side format, with each policy identified by name in the title bar on the left and right sides of the comparison view. When comparing two policies other than the running configuration, the time of last modification and the last user to modify are displayed with the policy name. Differences between the two policies are highlighted:

- Blue indicates that the highlighted setting is different in the two policies, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one policy but not the other.

You can perform any of the actions in the following table.

Table 15-6 SSL Policy Comparison View Actions

To...	You can...
navigate individually through changes	click **Previous** or **Next** above the title bar. The double-arrow icon ( ) centered between the left and right sides moves, and the **Difference** number adjusts to identify which difference you are viewing.
generate a new policy comparison view	click **New Comparison.** The **Select Comparison** window appears. See Using the SSL Policy Comparison Report, page 15-11 for more information.
generate a policy comparison report	click **Comparison Report**. The policy comparison report creates a PDF document that lists only the differences between the two policies.

Using the SSL Policy Comparison Report

License: Any

An SSL policy comparison report is a record of all differences between two SSL policies or a policy and the currently applied policy identified by the policy comparison view, presented in PDF format. You can use this report to further examine the differences between two policy configurations and to save and disseminate your findings.

You can generate an SSL policy comparison report from the comparison view for any policies to which you have access. Remember to save any changes before you generate a policy report; only saved changes appear in the report.

The format of the policy comparison report is the same as the policy report with one exception: the policy report contains all configurations in the policy, and the policy comparison report lists only those configurations that differ between the policies. An SSL policy comparison report contains the sections described in Generating a Report of Current Traffic Decryption Settings, page 15-9.

Tip You can use a similar procedure to compare access control, network analysis, intrusion, file, system, or health policies.

To compare two SSL policies:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > SSL.**

The SSL Policy appears.

Step 2 Click **Compare Policies**.

The Select Comparison window appears.

Step 3 From the **Compare Against** drop-down list, select the type of comparison you want to make:

- To compare two different policies, select **Other Policy**.

  The page refreshes and the Policy A and Policy B drop-down lists appear.
- To compare another policy to the currently active policy, select **Running Configuration.**

  The page refreshes and the Target/Running Configuration A and Policy B drop-down lists appear.

Step 4 Depending on the comparison type you selected, you have the following choices:

- If you are comparing two different policies, select the policies you want to compare from the Policy A and Policy B drop-down lists.
- If you are comparing the running configuration to another policy, select the second policy from the Policy B drop-down list.

Step 5 Click **OK** to display the policy comparison view.

The comparison view appears.

Step 6 Optionally, click **Comparison Report** to generate the SSL policy comparison report.

The SSL policy comparison report appears. Depending on your browser settings, the report may appear in a pop-up window, or you may be prompted to save the report to your computer.



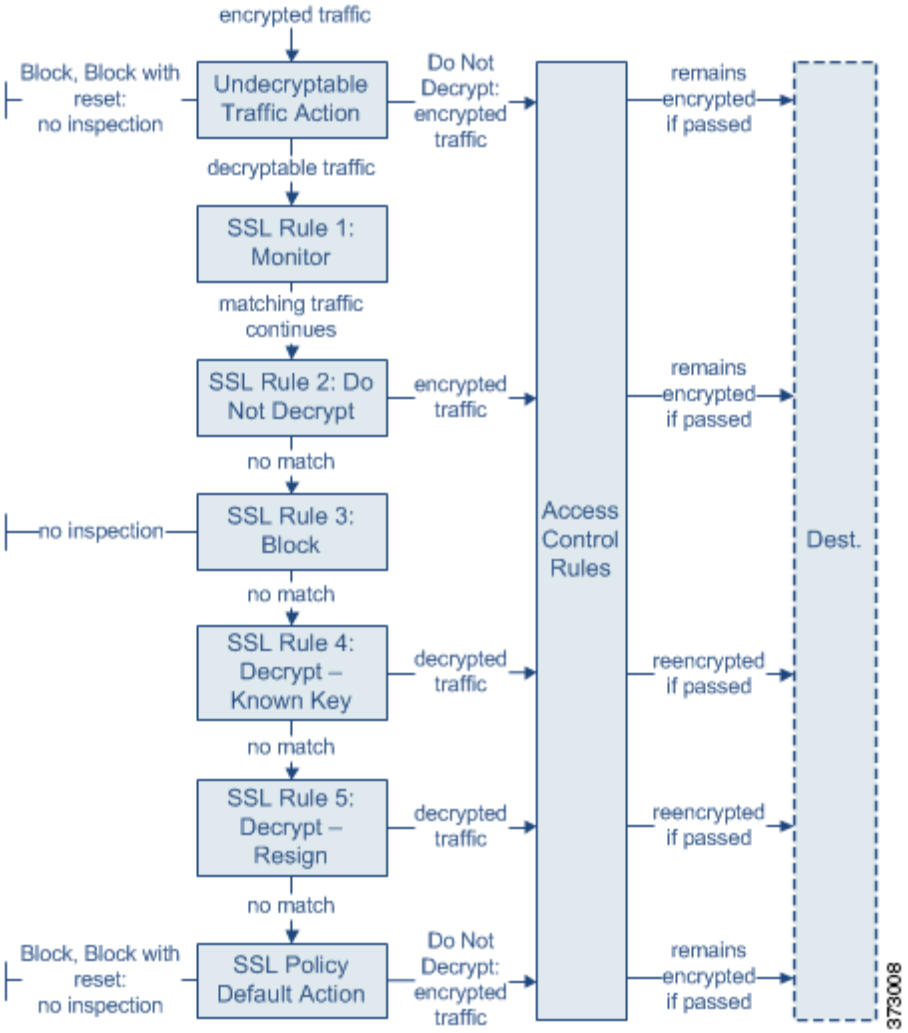
Getting Started with SSL Rules

Within an SSL policy, *SSL rules* provide a granular method of handling encrypted traffic, whether blocking the traffic without further inspection, not decrypting the traffic and inspecting it with access control, or decrypting the traffic for access control analysis.

The ASA FirePOWER module matches traffic to SSL rules in the order you specify. In most cases, the module handles encrypted traffic according to the *first* SSL rule where *all* the rule's conditions match the traffic. Conditions can be simple or complex; you can control traffic by security zone, network or geographical location, port, application, requested URL, user, certificate, certificate distinguished name, certificate status, cipher suite, or encryption protocol version.

Each rule also has an *action*, which determines whether you monitor, block, or inspect matching traffic with access control, optionally after decrypting matching traffic. Note that the module does **not** further inspect encrypted traffic it blocks. It does inspect encrypted and undecryptable traffic with access control. However, some access control rule conditions require unencrypted traffic, so encrypted traffic may match fewer rules. Also, by default, the module disables intrusion and file inspection of encrypted payloads.

The following scenario summarizes the ways that SSL rules handle traffic in an inline deployment.



In this scenario, traffic is evaluated as follows:

- **Undecryptable Traffic Action** evaluates encrypted traffic first. For traffic the module cannot decrypt, the module either blocks it without further inspection or passes it for access control inspection. Encrypted traffic that does not match continues to the next rule.
- **SSL Rule 1: Monitor** evaluates encrypted traffic next. Monitor rules track and log encrypted traffic but do not affect traffic flow. The module continues to match traffic against additional rules to determine whether to permit or deny it.
- **SSL Rule 2: Do Not Decrypt** evaluates encrypted traffic third. Matching traffic is not decrypted; the module inspects this traffic with access control, but not file or intrusion inspection. Traffic that does not match continues to the next rule.
- **SSL Rule 3: Block** evaluates encrypted traffic fourth. Matching traffic is blocked without further inspection. Traffic that does not match continues to the next rule.
- **SSL Rule 4: Decrypt - Known Key** evaluates encrypted traffic fifth. Matching traffic incoming to your network is decrypted using a private key you upload. The decrypted traffic is then evaluated against access control rules. Access control rules handle decrypted and unencrypted traffic

identically. The module can block traffic as a result of this additional inspection. All remaining traffic is reencrypted before being allowed to the destination. Traffic that does not match the SSL rule continues to the next rule.

- **SSL Rule 5: Decrypt - Resign** is the final rule. If traffic matches this rule, the module re-signs the server certificate with an uploaded CA certificate, then acts as a man-in-the-middle to decrypt traffic. The decrypted traffic is then evaluated against access control rules. Access control rules treat decrypted and unencrypted traffic identically. The module can block traffic as a result of this additional inspection. All remaining traffic is reencrypted before being allowed to the destination. Traffic that does not match the SSL rule continues to the next rule.
- **SSL Policy Default Action** handles all traffic that does not match any of the SSL rules. The default action either blocks encrypted traffic without further inspection or does not decrypt it, passing it for access control inspection.

For more information, see the following sections:

- Configuring Supporting Inspection Information, page 16-3
- Understanding and Creating SSL Rules, page 16-4
- Managing SSL Rules in a Policy, page 16-12

Configuring Supporting Inspection Information

License: Any

You must create reusable public key infrastructure (PKI) objects to control encrypted traffic based on encrypted session characteristics and decrypt encrypted traffic. You can add this information on the fly when uploading trusted certificate authority (CA) certificates to the SSL policy and creating SSL rule conditions, creating the associated object in the process. However, configuring these objects ahead of time reduces the chance of improper object creation.

Decrypting Encrypted Traffic with Certificates and Paired Keys

The ASA FirePOWER module can decrypt incoming encrypted traffic if you configure an internal certificate object by uploading the server certificate and private key used to encrypt the session. If you reference that object in an SSL rule with an action of **Decrypt - Known Key** and traffic matches that rule, the module uses the uploaded private key to decrypt the session.

The module can also decrypt outgoing traffic if you configure an internal CA object by uploading a CA certificate and private key. If you reference that object in an SSL rule with an action of **Decrypt - Resign** and traffic matches that rule, the module re-signs the server certificate passed to the client browser, then acts as a man-in-the-middle to decrypt the session.

See the following for more information:

- Working with Internal Certificate Objects, page 2-41
- Working with Internal Certificate Authority Objects, page 2-35

Controlling Traffic Based on Encrypted Session Characteristics

The ASA FirePOWER module can control encrypted traffic based on the cipher suite or server certificate used to negotiate the session. You can configure one of several different reusable objects and reference the object in an SSL rule condition to match traffic. The following table describes the different types of reusable objects you can configure:

If you configure...	You can control encrypted traffic based on whether...
a cipher suite list containing one or more cipher suites	the cipher suite used to negotiate the encrypted session matches a cipher suite in the cipher suite list
a trusted CA object by uploading a CA certificate your organization trusts	the trusted CA trusts the server certificate used to encrypt the session, whether: • the CA issued the certificate directly • the CA issued a certificate to an intermediate CA that issued the server certificate
an external certificate object by uploading a server certificate	the server certificate used to encrypt the session matches the uploaded server certificate
a distinguished name object containing a certificate subject or issuer distinguished name	the subject or issuer common name, country, organization, or organizational unit on the certificate used to encrypt the session matches the configured distinguished name

See the following for more information:

- Working with Geolocation Objects, page 2-42
- Working with Trusted Certificate Authority Objects, page 2-39
- Working with External Certificate Objects, page 2-41
- Working with Distinguished Name Objects, page 2-33

Understanding and Creating SSL Rules

License: Any

Within an SSL policy, SSL rules provide a granular method of handling network traffic. In addition to its unique name, each SSL rule has the following basic components.

State

By default, rules are enabled. If you disable a rule, the module does not use it to evaluate network traffic, and stops generating warnings and errors for that rule.

Position

Rules in an SSL policy are numbered, starting at 1. The module matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Conditions

Conditions specify the specific traffic the rule handles. Conditions can match traffic by security zone, network or geographical location, port, application, requested URL, user, certificate, certificate subject or issuer, certificate status, cipher suite, or encryption protocol version. Conditions can be simple or complex; their use can depends on device licenses.

Action

A rule's action determines how the module handles matching traffic. You can monitor, trust, block, or decrypt matching traffic. Decrypted traffic is subject to further inspection. Note that the module does **not** perform inspection on blocked or trusted encrypted traffic.

Logging

A rule's logging settings govern the records the module keeps of the traffic it handles. You can keep a record of traffic that matches a rule. You can log a connection when the module blocks an encrypted session or allows it to pass uninspected, according to the settings in an SSL policy. You can also force the module to log connections that it decrypts for further evaluation by access control rules, regardless of how the module later handles or inspects the traffic. You can log connections to the module log (syslog) or to an SNMP trap server.

Tip Properly creating and ordering SSL rules is a complex task, but one that is essential to building an effective deployment. If you do not plan your policy carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. To help ensure that the module handles traffic as you expect, the SSL policy interface has a robust warning and error feedback system for rules. For more information, see Troubleshooting SSL Rules, page 16-15.

To create or modify an SSL rule:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > SSL**.

The SSL Policy page appears.

Step 2 Click the edit icon ( ) next to the SSL policy where you want to add a rule.

The SSL policy editor appears, focused on the Rules tab.

Step 3 You have the following options:

- To add a new rule, click **Add Rule**.
- To edit an existing rule, click the edit icon ( ) next to the rule you want to edit.

The SSL rule editor appears.

Step 4 Type a **Name** for the rule.

Each rule must have a unique name. You can use up to thirty printable characters, including spaces and special characters, with the exception of the colon (:).

Step 5 Configure the rule components, as summarized above. You can configure the following, or accept the defaults:

- Specify whether the rule is **Enabled**.
- Specify the rule position; see Specifying an SSL Rule's Order of Evaluation, page 16-6.
- Select a rule **Action**; see Using Rule Actions to Determine Encrypted Traffic Handling and Inspection, page 16-8.
- Configure the rule's conditions; see Using Conditions to Specify the Encrypted Traffic a Rule Handles, page 16-6.
- Specify **Logging** options; see Logging Decryptable Connections with SSL Rules, page 36-14.

Step 6 Click **Save** to save the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Specifying an SSL Rule's Order of Evaluation

License: Any

When you first create an SSL rule, you specify its position using the **Insert** drop-down list in the rule editor. SSL rules in an SSL policy are numbered, starting at 1. The ASA FirePOWER module matches traffic to SSL rules in top-down order by ascending rule number.

In most cases, the module handles network traffic according to the *first* SSL rule where *all* the rule's conditions match the traffic. Except in the case of Monitor rules (which log traffic but do not affect traffic flow), the module does **not** continue to evaluate traffic against additional, lower-priority rules after that traffic matches a rule.

Tip Proper SSL rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs. For more information, see Ordering SSL Rules to Improve Performance and Avoid Preemption, page 16-16.

In addition to ordering rules by number, you can group rules by category. By default the module provides three categories: Administrator, Standard, and Root. You can add custom categories, but you cannot delete the ASA FirePOWER module-provided categories or change their order. For information on changing the position or category of an existing rule, see Changing an SSL Rule's Position or Category, page 16-13.

To add a rule to a category while editing or creating a rule:

Step 1 In the SSL rule editor, from the **Insert** drop-down list, select **Into Category**, then select the category you want to use.

When you save the rule, it is placed last in that category.

To position a rule by number while editing or creating a rule:

Step 1 In the SSL rule editor, from the **Insert** drop-down list, select **above rule** or **below rule**, then type the appropriate rule number.

When you save the rule, it is placed where you specified.

Using Conditions to Specify the Encrypted Traffic a Rule Handles

License: feature dependent

An SSL rule's conditions identify the type of encrypted traffic that rule handles. Conditions can be simple or complex, and you can specify more than one condition type per rule. Only if traffic meets all the conditions in a rule does the rule apply to the traffic.

If you do not configure a particular condition for a rule, the module does not match traffic based on that criterion. For example, a rule with a certificate condition but no version condition evaluates traffic based on the server certificate used to negotiate the session, regardless of the session SSL or TLS version.

When you add or edit an SSL rule, use the tabs on the left side of the lower portion of the rule editor to add and edit rule conditions. The conditions you can add to an SSL rule are described in the following table.

Table 16-1 SSL Rule Condition Types

This Condition...	Matches Encrypted Traffic...	Details
Zones	entering or leaving a device via an interface in a specific security zone	A security zone is a logical grouping of one or more interfaces according to your deployment and security policies. To build a zone condition, see Controlling Encrypted Traffic by Network Zone, page 17-2.
Networks	by its source or destination IP address, country, or continent	You can explicitly specify IP addresses. The geolocation feature also allows you to control traffic based on its source or destination country or continent. To build a network condition, see Controlling Encrypted Traffic by Network or Geographical Location, page 17-3.
Ports	by its source or destination port	You can control encrypted traffic based on the TCP port. To build a port condition, see Controlling Encrypted Traffic by Port, page 17-5.
Users	by the user involved in the session	You can control encrypted traffic based on the LDAP user logged into a host involved in an encrypted, monitored session. You can control traffic based on individual users or groups retrieved from a Microsoft Active Directory server. To build a user condition, see Controlling Encrypted Traffic Based on User, page 17-6.
Applications	by the application detected in a session	You can control access to individual applications in encrypted sessions, or filter access according to basic characteristics: type, risk, business relevance, and categories. To build an application condition, see Controlling Encrypted Traffic Based on Application, page 17-8.
Categories	by the URL requested in the session, based on the certificate subject distinguished name	You can limit the websites that users on your network can access based on the URL's general classification and risk level. To build a URL condition, see Controlling Encrypted Traffic by URL Category and Reputation, page 17-13.
Distinguished Names	by the subject or issuer distinguished name of the server certificate used to negotiate the encrypted session	You can control encrypted traffic based on the CA that issued a server certificate, or the server certificate holder. To build a distinguished name condition, see Controlling Encrypted Traffic by Certificate Distinguished Name, page 17-17.
Certificates	by the server certificate used to negotiate the encrypted session	You can control encrypted traffic based on the server certificate passed to the user's browser in order to negotiate the encrypted session. To build a certificate condition, see Controlling Encrypted Traffic by Certificate Status, page 17-20.

Table 16-1 SSL Rule Condition Types (continued)

This Condition...	Matches Encrypted Traffic...	Details
Certificate Status	by properties of the server certificate used to negotiate the encrypted session	You can control encrypted traffic based on a server certificate's status. To build a certificate status condition, see Controlling Encrypted Traffic by Certificate Status, page 17-20.
Cipher Suites	by the cipher suite used to negotiate the encrypted session	You can control encrypted traffic based on the cipher suite selected by the server to negotiate the encrypted session. To build a cipher suite condition, see Controlling Encrypted Traffic by Cipher Suite, page 17-25.
Versions	by the version of SSL or TLS used to encrypt the session	You can control encrypted traffic based on the version of SSL or TLS used to encrypt the session. To build a version condition, see Controlling Traffic by Encryption Protocol Version, page 17-26.

Note that while you can control and inspect encrypted traffic, controlling traffic using detected application, URL category, or user requires additional licenses. Also, overly complex rules can consume excessive resources and in some cases prevent you from applying the policy. For more information, see Troubleshooting SSL Rules, page 16-15.

Using Rule Actions to Determine Encrypted Traffic Handling and Inspection

License: Any

Every SSL rule has an associated action that determines the following for matching encrypted traffic:

- handling — foremost, the rule action governs whether the ASA FirePOWER module will monitor, trust, block, or decrypt encrypted traffic that matches the rule's conditions
- logging — the rule action determines when and how you can log details about matching encrypted traffic.

Your SSL inspection configuration handles, inspects, and logs decrypted traffic:

- The SSL policy's undecryptable actions handle traffic that the ASA FirePOWER module cannot decrypt; see Setting Default Handling for Undecryptable Traffic, page 15-4.
- The policy's default action handles traffic that does not meet the condition of any non-Monitor SSL rule; see Setting Default Handling and Inspection for Encrypted Traffic, page 15-3.

You can log a connection event when the ASA FirePOWER module blocks or trusts an encrypted session. You can also force the module to log connections that it decrypts for further evaluation by access control rules, regardless of how the module later handles or inspects the traffic. Connection logs for encrypted sessions contain details about the encryption, such as the certificate used to encrypt that session. You can log only end-of-connection events, however:

- for blocked connections (Block, Block with reset), the module immediately ends the sessions and generates an event
- for trusted connections (Do not decrypt), the module generates an event when the session ends

For detailed information on rule actions and how they affect handling and logging, see the following sections:

- Monitor Action: Postponing Action and Ensuring Logging, page 16-9

- Do Not Decrypt Action: Passing Encrypted Traffic Without Inspection, page 16-9
- Blocking Actions: Blocking Encrypted Traffic Without Inspection, page 16-9
- Decrypt Actions: Decrypting Traffic for Further Inspection, page 16-10
- Managing SSL Rules in a Policy, page 16-12

Monitor Action: Postponing Action and Ensuring Logging

License: Any

The **Monitor** action does not affect encrypted traffic flow; matching traffic is neither immediately permitted nor denied. Rather, traffic is matched against additional rules, if present, to determine whether to trust, block, or decrypt it. The first non-Monitor rule matched determines traffic flow and any further inspection. If there are no additional matching rules, the ASA FirePOWER module uses the default action.

Because the primary purpose of Monitor rules is to track network traffic, the module automatically logs end-of connection events for monitored traffic. That is, the module always logs the end of the connection, regardless of the logging configuration of the rule or default action that later handles the connection. In other words, if a packet matches a Monitor rule, the connection is always logged, even if the packet matches no other rules and you do not enable logging on the default action.

Do Not Decrypt Action: Passing Encrypted Traffic Without Inspection

License: Any

The **Do not decrypt** action passes encrypted traffic for evaluation by the access control policy's rules and default action. Because some access control rule conditions require unencrypted traffic, this traffic may match fewer rules. The module cannot perform deep inspection on encrypted traffic, such as intrusion or file inspection.

Blocking Actions: Blocking Encrypted Traffic Without Inspection

License: Any

The **Block** and **Block with reset** actions are analogous to the access control rule actions Block and Block with reset. These actions prevent the client and server from establishing the SSL-encrypted session and passing encrypted traffic. Block with reset rules also reset the connection.

Note that the ASA FirePOWER module does not display the configured response page for blocked encrypted traffic. Instead, users requesting prohibited URLs have their connection either reset or time out. See Displaying a Custom Web Page for Blocked URLs, page 8-14 for more information.

Tip Note that you cannot use the Block or Block with reset action in a passive or inline (tap mode) deployment, as the device does not directly inspect the traffic. If you create a rule with the Block or Block with reset action that contains passive or inline (tap mode) interfaces within a security zone condition, the policy editor displays a warning icon ( ) next to the rule.

Decrypt Actions: Decrypting Traffic for Further Inspection

License: Any

The **Decrypt - Known Key** and **Decrypt - Resign** actions decrypt encrypted traffic. The ASA FirePOWER module inspects decrypted traffic with access control. Access control rules handle decrypted and unencrypted traffic identically — you can detect and block intrusions, prohibited files, and malware. The module reencrypts allowed traffic before passing it to its destination.

When you configure the **Decrypt - Known Key** action, you can associate one or more server certificates and paired private keys with the action. If traffic matches the rule, and the certificate used to encrypt the traffic matches the certificate associated with the action, the module uses the appropriate private key to obtain the session encryption and decryption keys. Because you must have access to the private key, this action is best suited to decrypt traffic incoming to servers your organization controls.

Similarly, you can associate one Certificate Authority certificate and private key with the **Decrypt - Resign** action. If traffic matches this rule, the module re-signs the server certificate with the CA certificate, then acts as a man-in-the-middle. It creates two SSL sessions, one between client and device, one between device and server. Each session contains different cryptographic session details, and allows the module to decrypt and reencrypt traffic. This action is more suited for outgoing traffic, as you replace the certificate's private key with one you control to obtain the session keys.

Re-signing a server certificate involves either replacing the certificate's public key with a CA certificate public key, or replacing the entire certificate. Normally, if you replace an entire server certificate, the client browser warns the certificate is not signed by a trusted authority when establishing the SSL connection. However, if your client's browser trusts the CA in the policy, the browser does not warn that the certificate is not trusted. If the original server certificate is self-signed, the ASA FirePOWER module replaces the entire certificate, and trusts the re-signing CA, but the user's browser does not warn that the certificate is self-signed. In this case, replacing only the server certificate public key causes the client browser does warn that the certficate is self-signed.

If you configure a rule with the **Decrypt - Resign** action, the rule matches traffic based on the referenced internal CA certificate's signature algorithm type, in addition to any configured rule conditions. Because you associate one CA certificate with a **Decrypt - Resign** action, you cannot create an SSL rule that decrypts multiple types of outgoing traffic encrypted with different signature algorithms. In addition, any external certificate objects and cipher suites you add to the rule must match the associated CA certificate encryption algorithm type.

For example, outgoing traffic encrypted with an elliptic curve (EC) algorithm matches a **Decrypt - Resign** rule only if the action references an EC-based CA certificate; you must add EC-based external certificates and cipher suites to the rule if you want to create certificate and cipher suite rule conditions. Similarly, a **Decrypt - Resign** rule that references an RSA-based CA certificate matches only outgoing traffic encrypted with an RSA algorithm; outgoing traffic encrypted with an EC algorithm does not match the rule, even if all other configured rule conditions match.

Note the following:

- You cannot use the **Decrypt - Known Key** action in a passive deployment if the cipher suite used to establish the SSL connection applies either the Diffie-Hellman ephemeral (DHE) or the elliptic curve Diffie-Hellman ephemeral (ECDHE) key exchange algorithm. If your SSL policy targets passive or inline (tap mode) interfaces, and contains a **Decrypt - Known Key** rule with a cipher suite condition containing either a DHE or an ECDHE cipher suite, the ASA FirePOWER module displays an information icon ( ) next to the rule. If you later add a zone condition to the SSL rule that contains passive or inline (tap mode) interfaces, the module displays a warning icon ( ).
- You cannot use the **Decrypt - Resign** action in a passive or inline (tap mode) deployment, as the device does not directly inspect traffic. If you create a rule with the **Decrypt - Resign** action that contains passive or inline (tap mode) interfaces within a security zone, the policy editor displays a warning

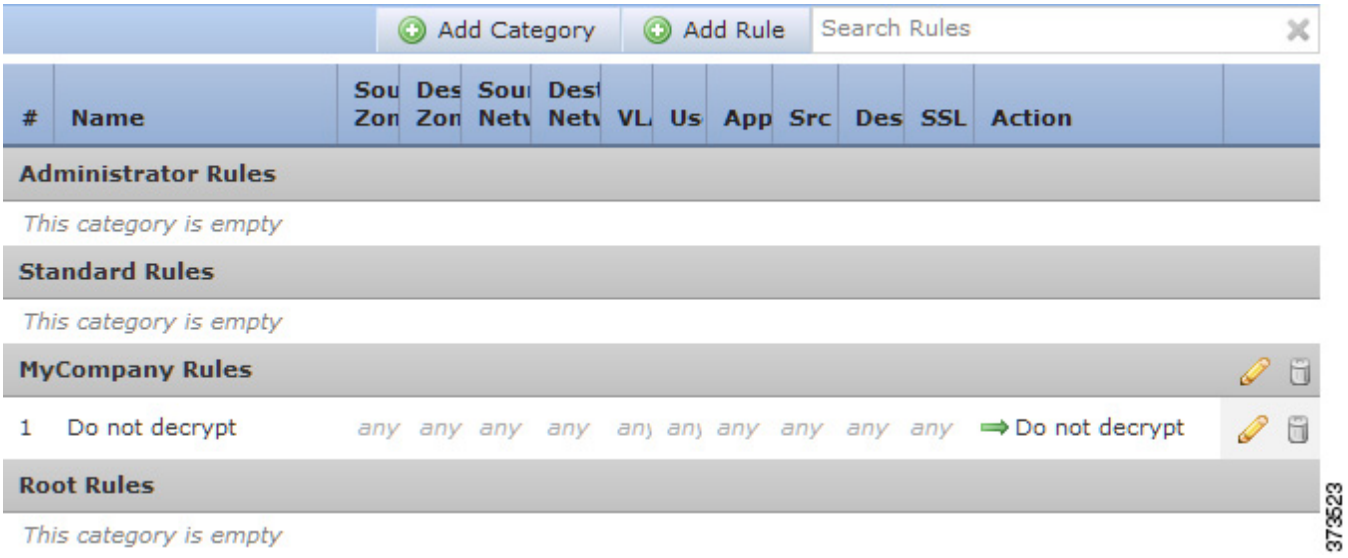
icon ( ) next to the rule. If your SSL policy targets passive or inline (tap mode) interfaces, and contains a **Decrypt - Resign** rule, the module displays an information icon ( ) next to the rule. If you later add a zone condition to the SSL rule that contains passive or inline (tap mode) interfaces, the module displays a warning icon ( ). If you apply an SSL policy that contains a Decrypt - Resign rule to a device with passive or inline (tap mode) interfaces, any SSL sessions that match the rule fail.

- If the client does not trust the CA used to re-sign the server certificate, it warns the user that the certificate should not be trusted. To prevent this, import the CA certificate into the client trusted CA store. Alternatively, if your organization has a private PKI, you can issue an intermediate CA certificate signed by the root CA which is automatically trusted by all clients in the organization, then upload that CA certificate to the device.
- You can add an anonymous cipher suite to the **Cipher Suite** condition in an SSL rule, but keep in mind:
  - The system automatically strips anonymous cipher suites during ClientHello processing. For the system to use the rule, you must also configure your SSL rules in an order that prevents ClientHello processing. For more information, see Ordering SSL Rules to Improve Performance and Avoid Preemption, page 16-16.
  - You cannot use the **Decrypt - Resign** or **Decrypt - Known Key** action in the rule, because the system cannot decrypt traffic encrypted with an anonymous cipher suite.
- The ASA FirePOWER module cannot decrypt traffic if an HTTP proxy is positioned between a client and your device, and the client and server establish a tunneled SSL connection using the CONNECT HTTP method. The **Handshake Errors** undecryptable action determines how the module handles this traffic. See Setting Default Handling for Undecryptable Traffic, page 15-4 for more information.
- You cannot match on **Distinguished Name** or **Certificate** conditions when creating an SSL rule with a **Decrypt - Known Key** action. The assumption is that if this rule matches traffic, the certificate, subject DN, and issuer DN already match the certificate associated with the rule. For more information, see Using Rule Actions to Determine Encrypted Traffic Handling and Inspection, page 16-8.
- If you create an internal CA object and choose to generate a certificate signing request (CSR), you cannot use this CA for a **Decrypt - Resign** action until you upload the signed certificate to the object. For more information, see Obtaining and Uploading a New Signed Certificate, page 2-37.
- If you configure a rule with the **Decrypt - Resign** action, and mismatch signature algorithm type for one or more external certificate objects or cipher suites, the policy editor displays an information icon ( ) next to the rule. If you mismatch signature algorithm type for all external certificate objects, or all cipher suites, the policy displays a warning icon ( ) next to the rule, and you cannot apply the access control policy associated with the SSL policy. For more information, see Controlling Encrypted Traffic by Certificate, page 17-19 and Controlling Encrypted Traffic by Cipher Suite, page 17-25.
- If decrypted traffic matches an access control rule with an action of **Interactive Block** or **Interactive Block with reset**, the ASA FirePOWER module blocks the matching connection without interaction and the module does **not** display a response page.
- If you enable the **Normalize Excess Payload** option in the inline normalization preprocessor, when the preprocessor normalizes decrypted traffic, it may drop a packet and replace it with a trimmed packet. This does not end the SSL session. If the traffic is allowed, the trimmed packet is encrypted as part of the SSL session. For more information on this option, see Normalizing Inline Traffic, page 24-6.
- If your browser uses certificate pinning to verify a server certificate, you cannot decrypt this traffic by re-signing the server certificate. If you want to allow this traffic, configure an SSL rule with the Do not decrypt action to match the server certificate common name or distinguished name.

Managing SSL Rules in a Policy

License: Any

The Rules tab of the SSL policy editor, shown in the following graphic, allows you to add, edit, search, move, enable, disable, delete, and otherwise manage SSL rules within your policy.



For each rule, the policy editor displays its name, a summary of its conditions, and the rule action. Icons represent warnings, errors, and other important information. Disabled rules are grayed out and marked `(disabled)` beneath the rule name. See Troubleshooting SSL Rules, page 16-15 for more information about the icons.

For information on managing SSL rules, see:

- Searching SSL Rules, page 16-12
- Enabling and Disabling SSL Rules, page 16-13
- Changing an SSL Rule's Position or Category, page 16-13

Searching SSL Rules

License: Any

You can search the list of SSL rules for matching values using an alphanumeric string, including spaces and printable, special characters. The search inspects the rule name and any rule condition you have added to the rule. For rule conditions, the search matches any name or value you can add for each condition type (zone, network, application, and so on). This includes individual object names or values, group object names, individual object names or values within a group, and literal values.

You can use complete or partial search strings. The column for matching values is highlighted for each matching rule. For example, if you search on all or part of the string `100Bao`, at a minimum, the Applications column is highlighted for each rule where you have added the 100Bao application. If you also have a rule named `100Bao`, both the Name and Applications columns are highlighted.

You can navigate to each previous or next matching rule. A status message displays the current match and the total number of matches.

Matches may occur on any page of a multi-page rule list. When the first match is not on the first page, the page where the first match occurs is displayed. Selecting the next match when you are at the last match takes you to the first match, and selecting the previous match when you are at the first match takes you to the last match.

To search for rules:

Step 1 In the SSL policy editor for the policy you want to search, click the **Search Rules** prompt, type a search string, then press Enter. You can also use the Tab key or click a blank page area to initiate the search.

Columns for rules with matching values are highlighted, with differentiated highlighting for the indicated (first) match.

Step 2 Find the rules you are interested in:

- To navigate between matching rules, click the next-match (▼) or previous-match (▲) icon.
- To refresh the page and clear the search string and any highlighting, click the clear icon (✖).

Enabling and Disabling SSL Rules

License: Any

When you create an SSL rule, it is enabled by default. If you disable a rule, the ASA FirePOWER module does not use it to evaluate network traffic and stops generating warnings and errors for that rule. When viewing the list of rules in an SSL policy, disabled rules are grayed out, although you can still modify them. Note that you can also enable or disable an SSL rule using the rule editor; see Understanding and Creating SSL Rules, page 16-4.

To change an SSL rule's state:

Step 1 In the SSL policy editor for the policy that contains the rule you want to enable or disable, right-click the rule and choose a rule state:

- To enable an inactive rule, select **State > Enable**.
- To disable an active rule, **State > Disable**.

Step 2 Click **Store ASA FirePOWER Changes**.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Changing an SSL Rule's Position or Category

License: Any

To help you organize SSL rules, every SSL policy has three ASA FirePOWER module-provided rule categories: Administrator Rules, Standard Rules, and Root Rules. You cannot move, delete, or rename these categories, although you can create custom categories.

For more information, see:

- Moving an SSL Rule, page 16-14
- Adding a New SSL Rule Category, page 16-14

Moving an SSL Rule

License: Any

Proper SSL rule order reduces the resources required to process network traffic, and prevents rule preemption.

The following procedure explains how to move one or more rules at a time using the SSL policy editor. You can also move individual SSL rules using the rule editor; see Understanding and Creating SSL Rules, page 16-4.

To move a rule:

Step 1 In the SSL policy editor for the policy that contains the rules you want to move, select the rules by clicking in a blank area for each rule. Use the Ctrl and Shift keys to select multiple rules.

The rules you selected are highlighted.

Step 2 Move the rules. You can cut and paste or drag and drop.

To cut and paste rules into a new location, right-click a selected rule and select **Cut**. Then, right-click a blank area for a rule next to where you want to paste the cut rules and select **Paste above** or **Paste below**. Note that you cannot copy and paste SSL rules between two different SSL policies.

Step 3 Click **Store ASA FirePOWER Changes**..

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Adding a New SSL Rule Category

License: Any

To help you organize SSL rules, every SSL policy has three ASA FirePOWER module-provided rule categories: Administrator Rules, Standard Rules, and Root Rules. You cannot move, delete, or rename these categories, although you can create custom categories between the Standard Rules and Root Rules.

Adding custom categories allows you to further organize your rules without having to create additional policies. You can rename and delete categories that you add. You cannot move these categories, but you can move rules into, within, and out of them.

To add a new category:

Step 1 In the SSL policy editor for the policy where you want to add a rule category, click **Add Category**.

Tip If your policy already contains rules, you can click a blank area in the row for an existing rule to set the position of the new category before you add it. You can also right-click an existing rule and select **Insert new category**.

The Add Category pop-up window appears.

Step 2 Type a unique category **Name**.

You can enter an alphanumeric name, including spaces and special printable characters, with up to 30 characters.

Step 3 You have the following choices:

- To position the new category immediately above an existing category, select **above Category** from the first **Insert** drop-down list, then select the category above which you want to position the rule from the second drop-down list.
- To position the new category rule below an existing rule, select **below rule** from the drop-down list, then enter an existing rule number. This option is valid only when at least one rule exists in the policy.
- To position the rule above an existing rule, select **above rule** from the drop-down list, then, enter an existing rule number. This option is valid only when at least one rule exists in the policy.

Step 4 Click **OK**.

Your category is added. You can click the edit icon ( ) next to a custom category to edit its name, or click the delete icon ( ) to delete the category. Rules in a category you delete are added to the category above.

Step 5 Click **Store ASA FirePOWER Changes** to save the policy.

Troubleshooting SSL Rules

License: Any

Properly creating and ordering SSL rules is a complex task, but one that is essential to building an effective deployment. If you do not plan your policy carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. To help ensure that the ASA FirePOWER module handles traffic as you expect, the SSL policy interface has a robust warning and error feedback system for rules.

For each rule, icons in the policy editor mark warnings and errors, as described in the following table. Hover your pointer over the icon to read the warning, error, or informational text.

Table 16-2 SSL Error Icons

Icon	Description	Details
	warning	Depending on the issue, you may be able to apply an SSL policy that displays rule or other warnings. In these cases, the misconfigured settings will have no effect. For example, a preempted rule never evaluates traffic. However, if a warning icon marks a licensing error or model mismatch, you cannot apply the policy until you correct the issue. If you disable a rule with a warning, the warning icon disappears. It reappears if you enable the rule without correcting the underlying issue.

Table 16-2 SSL Error Icons (continued)

Icon	Description	Details
	error	If a rule or other SSL policy configuration has an error, you cannot apply the policy until you correct the issue.
	information	Information icons convey helpful information about configurations that may affect the flow of traffic. These issues are minor and will not prevent you from applying the policy.

Properly configuring SSL rules can also reduce the resources required to process network traffic. Creating complex rules and mis-ordering rules can affect performance.

For more information, see:

- Understanding Rule Preemption and Invalid Configuration Warnings, page 16-16
- Ordering SSL Rules to Improve Performance and Avoid Preemption, page 16-16

Understanding Rule Preemption and Invalid Configuration Warnings

License: Any

Properly configuring and ordering SSL rules is essential to building an effective deployment. Within an SSL policy, SSL rules can preempt other rules or contain invalid configurations. The module uses warning and error icons to mark these issues.

Understanding Rule Preemption Warnings

The conditions of an SSL rule may preempt a subsequent rule from matching traffic. For example:

```
Rule 1: do not decrypt Administrators
Rule 2: block Administrators
```

The second rule above will never block traffic because the first rule will have already allowed the traffic.

Understanding Invalid Configuration Warnings

Because outside settings that the SSL policy depends on may change, an SSL policy setting that was valid may become invalid. Consider the following examples:

- A rule that contains a URL category condition might be valid until you target a module that does not have a URL Filtering license. At that point, an error icon appears next to the rule, and you cannot apply the policy to that device until you edit or delete the rule, retarget the policy, or enable the appropriate license.
- If you create a Decrypt - Resign rule, and later add a security zone with passive interfaces to a zone condition, the module displays a warning icon next to the rule. Because you cannot decrypt traffic by re-signing a certificate in a passive deployment, the rule has no effect until you remove the passive interfaces from the rule or change the rule action.
- If you add a user to a rule, then change your LDAP user awareness settings to exclude that user, the rule will have no effect because the user is no longer an access-controlled user.

Ordering SSL Rules to Improve Performance and Avoid Preemption

License: Any

Rules in an SSL policy are numbered, starting at 1. The ASA FirePOWER module matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Proper SSL rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs.

Order Rules from Most to Least Critical

First, you must order rules to suit your organization's needs. Place priority rules that must apply to all traffic near the top of the policy. For example, if you want to decrypt outgoing traffic from a single user for further analysis (using a Decrypt - Resign rule), but not decrypt traffic from all other users in the department (using a Do not decrypt rule), place two SSL rules in that order.

Order Rules from Specific to General

You can improve performance by placing specific rules earlier, that is, rules that narrowly define the traffic they handle. This is also important because rules with broad conditions can match many different types of traffic, and can preempt later, more specific rules.

Consider a scenario where a trusted CA (Good CA) mistakenly issued a CA certificate to a malicious entity (Bad CA), but has not yet revoked that certificate. You want to block traffic encrypted with certificates issued by the untrusted CA, but otherwise allow traffic within the trusted CA's chain of trust. You should upload the CA certificates and all intermediate CA certificates, then order your rules as follows:

```
Rule 1: Block issuer CN=www.badca.com
Rule 2: Do not decrypt issuer CN=www.goodca.com
```

If you reverse the rules:

```
Rule 1: Do not decrypt issuer CN=www.goodca.com
Rule 2: Block issuer CN=www.badca.com
```

the first rule matches all traffic trusted by Good CA, including traffic trusted by Bad CA. Because no traffic ever matches the second rule, malicious traffic may be allowed instead of blocked.

Order Rules to Allow Traffic from Certificate Pinned Sites

Certificate pinning forces a client's browser to verify that a server's public key certificate matches a certificate the browser already associated with the server before establishing an SSL session. Because the Decrypt - Resign action involves modifying a server certificate before passing it to the client, these modified certificates are rejected if the browser already pinned that certificate.

For example, if a client browser connects to `windowsupdate.microsoft.com`, a site that uses certificate pinning, and you configure an SSL rule that matches that traffic with a Decrypt - Resign action, the ASA FirePOWER module re-signs the server certificate before passing it to the client browser. Because this modified server certificate does not match the browser's pinned certificate for `windowsupdate.microsoft.com`, the client browser rejects the connection.

If you want to allow this traffic, configure an SSL rule with the Do not decrypt action to match the server certificate common name or distinguished name. In the SSL policy, order this rule before all Decrypt - Resign rules that also match the traffic. You can retrieve the pinned certificate from the client's browser after a successful connection to the website. You can also view the certificate from the logged connection event, whether the connection succeeded or failed.

Place Rules that Decrypt Traffic Later

Because traffic decryption requires processing resources, placing rules that do not decrypt traffic (Do not decrypt, Block) before rules that do (Decrypt - Known Key, Decrypt - Resign) can improve performance. This is because traffic decryption can command significant resources. In addition, Block rules can divert traffic that the ASA FirePOWER module might otherwise have decrypted or inspected. All other factors being equal, that is, given a set of rules where none is more critical and preemption is not an issue, consider placing them in the following order:

- Monitor rules that log matching connections, but take no other action on traffic
- Block rules that block traffic without further inspection
- Do not decrypt rules that do not decrypt encrypted traffic
- Decrypt - Known Key rules that decrypt incoming traffic with a known private key
- Decrypt - Resign rules that decrypt outgoing traffic by re-signing the server certificate

Prioritize ClientHello Modifications

To prioritize ClientHello modifications, place rules that match on conditions that are available in the ClientHello message before rules that match on ServerHello or server Certificate conditions.

When a managed device processes an SSL handshake, it can modify the ClientHello message to increase the likelihood of decryption. For example, it may remove compression methods because the Firepower System cannot decrypt compressed sessions.

The system only modifies ClientHello messages if it can conclusively match them to an SSL rule with a Decrypt - Resign action. The first time the system detects an encrypted session to a new server, server Certificate data is not available for ClientHello processing, which can result in an undecrypted first session. For subsequent connections from the same client, the system can match the ClientHello message conclusively to rules with server Certificate conditions and process the message to maximize decryption potential.

If you place rules that match on ServerHello or server Certificate conditions (certificate, distinguished names, certificate status, cipher suites, version) before rules that match on ClientHello conditions (zones, networks, VLAN tags, ports, users, applications, URL categories), you can preempt ClientHello modification and increase the number of undecrypted sessions.

Configuring SSL Inspection to Improve Performance

License: Any

Complex SSL policies and rules can command significant resources. When you apply an SSL policy, the ASA FirePOWER module evaluates all the rules together and creates an expanded set of criteria that the device uses to evaluate network traffic. A pop-up window may warn that you have exceeded the maximum number of SSL rules supported by a device. This maximum depends on a number of factors, including the physical memory and the number of processors on the device.

Simplifying Rules

The following guidelines can help you simplify your SSL rules and improve performance:

- When constructing a rule, use as few individual elements in your conditions as possible. For example, in network conditions, use IP address blocks rather than individual IP addresses. In port conditions, use port ranges. Use application filters and URL categories and reputations to perform application control and URL filtering, and LDAP user groups to perform user control.

Note that combining elements into objects that you then use in SSL rule conditions does not improve performance. For example, using a network object that contains 50 individual IP addresses gives you only an organizational—not a performance—benefit over including those IP addresses in the condition individually.

- Restrict rules by security zones whenever possible. If a device's interfaces are not in one of the zones in a zone-restricted rule, the rule does not affect performance on that device.
- Do not overconfigure rules. If one condition is enough to match the traffic you want to handle, do not use two.

Configuring Traffic Decryption

Keep the following guidelines in mind when configuring traffic decryption:

- Traffic decryption requires processing resources to decrypt the traffic, and to inspect it with access control. Create narrowly focused decrypt rules over broad decrypt rules to reduce the amount of traffic the ASA FirePOWER module decrypts, and as a result, reduce the processing resources required to decrypt traffic. Rather than decrypting then later allowing or blocking traffic using an access control rule, block or choose not to decrypt encrypted traffic where possible.
- If you configure certificate status conditions to trust traffic based on the root issuer CA, upload the root CA certificate and all intermediate CA certificates within the root CA's chain of trust to your SSL policy. All traffic within a trusted CA's chain of trust can be allowed without decryption, rather than unnecessarily decrypting it.



Tuning Traffic Decryption Using SSL Rules

A basic SSL rule applies its rule action to all encrypted traffic inspected by the ASA FirePOWER module. To better control and decrypt encrypted traffic, you can configure rule conditions to handle and log specific types of traffic. Each SSL rule can contain 0, 1, or more rule conditions; a rule only matches traffic if the traffic matches every condition in that SSL rule.

Note When traffic matches a rule, the ASA FirePOWER module applies the configure rule action to the traffic. When the connection ends, the module logs the traffic if configured to do so. For more information, see Using Rule Actions to Determine Encrypted Traffic Handling and Inspection, page 16-8 and Logging Connections Based on Access Control Handling, page 36-9.

Each rule condition allows you to specify one or more properties of traffic you want to match against; these properties include details of:

- the flow of traffic, including the security zone through which it travels, IP address and port, and country of origin or destination
- the user associated with a detected IP address
- the traffic payload, including the application detected in the traffic
- the connection encryption, including the SSL/TLS protocol version and cipher suite and server certificate used to encrypt the connection
- the category and reputation of the URL specified in the server certificate's distinguished name

For more information, see the following sections:

- Logging Decryptable Connections with SSL Rules, page 36-14
- Controlling Encrypted Traffic with Network-Based Conditions, page 17-1
- Controlling Encrypted Traffic by Reputation, page 17-7
- Controlling Traffic Based on Server Certificate Characteristics, page 17-16

Controlling Encrypted Traffic with Network-Based Conditions

License: Any

SSL *rules* within *SSL policies* exert granular control over encrypted traffic logging and handling. Network-based conditions allow you to manage which encrypted traffic can traverse your network, using one or more of the following criteria:

- source and destination security zones
- source and destination IP addresses or geographical locations
- source and destination port

You can combine network-based conditions with each other and with other types of conditions to create an SSL rule. These SSL rules can be simple or complex, matching and inspecting traffic using multiple conditions. For detailed information on SSL rules, see Getting Started with SSL Rules, page 16-1.

For more information, see the following sections:

- Controlling Encrypted Traffic by Network Zone, page 17-2
- Controlling Encrypted Traffic by Network or Geographical Location, page 17-3
- Controlling Encrypted Traffic by Port, page 17-5

Controlling Encrypted Traffic by Network Zone

License: Any

Zone conditions in SSL rules allow you to control encrypted traffic by its source and destination security zones.

A *security zone* is a grouping of one or more interfaces. An option you choose during a device's initial setup, called its *detection mode*, determines how the ASA FirePOWER module initially configures the device's interfaces, and whether those interfaces belong to a security zone.

As a simple example, when you register a device with an **Inline** detection mode, the ASA FirePOWER module creates two zones: Internal and External, and assigns the first pair of interfaces on the device to those zones. Hosts connected to the network on the Internal side represent your protected assets.

Tip You are not required to group all internal (or external) interfaces into a single zone. Choose the grouping that makes sense for your deployment and security policies. For more information on creating zones, see Working with Security Zones, page 2-32.

In this deployment, you may decide that although you want these hosts to have unrestricted access to the Internet, you nevertheless want to protect them by decrypting and inspecting incoming encrypted traffic.

To accomplish this with SSL inspection, configure an SSL rule with a zone condition where the **Destination Zone** is set to **Internal**. This simple SSL rule matches traffic that leaves the device from any interface in the Internal zone.

If you want to build a more complex rule, you can add a maximum of 50 zones to each of the **Sources Zones** and **Destination Zones** in a single zone condition:

- To match encrypted traffic *leaving* the device from an interface in the zone, add that zone to the **Destination Zones**.

  Because devices deployed passively do not transmit traffic, you cannot use a zone comprised of passive interfaces in a **Destination Zone** condition.
- To match encrypted traffic *entering* the device from an interface in the zone, add that zone to the **Source Zones**.

If you add both source and destination zone conditions to a rule, matching traffic must originate from one of the specified source zones **and** egress through one of the destination zones.

Note that just as all interfaces in a zone must be of the same type (all inline, all passive, all switched, or all routed), all zones used in a zone condition for an SSL rule must be of the same type. That is, you cannot write a single rule that matches encrypted traffic to or from zones of different types.

Warning icons indicate invalid configurations, such as zones that contain no interfaces. For details, hover your pointer over the icon.

To control encrypted traffic by zone:

Step 1 In the SSL policy where you want to control encrypted traffic by zone, create a new SSL rule or edit an existing rule.

For detailed instructions, see Understanding and Creating SSL Rules, page 16-4.

Step 2 In the SSL rule editor, select the Zones tab.

The Zones tab appears.

Step 3 Find and select the zones you want to add from the **Available Zones**.

To search for zones to add, click the **Search by name** prompt above the **Available Zones** list, then type a zone name. The list updates as you type to display matching zones.

Click to select a zone. To select multiple zones, use the Shift and Ctrl keys, or right-click and then select **Select All**.

Step 4 Click **Add to Source** or **Add to Destination** to add the selected zones to the appropriate list.

You can also drag and drop selected zones.

Step 5 Save or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Controlling Encrypted Traffic by Network or Geographical Location

License: Any

Network conditions in SSL rules allow you to control and decrypt encrypted traffic by its source and destination IP address. You can either:

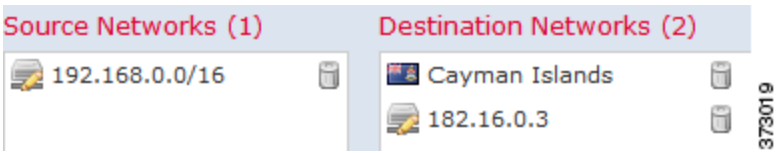
- explicitly specify the source and destination IP addresses for the encrypted traffic you want to control, or
- use the geolocation feature, which associates IP addresses with geographical locations, to control encrypted traffic based on its source or destination country or continent

When you build a network-based SSL rule condition, you can manually specify IP address and geographical locations. Alternately, you can configure network conditions with network and geolocation *objects*, which are reusable and associate a name with one or more IP addresses, address blocks, countries, continents, and so on.

Tip After you create a network or geolocation object, you can use it not only to build SSL rules, but also to represent IP addresses in various other places in the module interface. You can create these objects using the object manager; you can also create network objects on-the-fly while you are configuring SSL rules. For more information, see Managing Reusable Objects, page 2-1.

Note that if you want to write rules to control traffic by geographical location, to ensure you are using up-to-date geolocation data to filter your traffic, Cisco **strongly** recommends you regularly update the geolocation database (GeoDB) on your ASA FirePOWER module; see Updating the Geolocation Database, page 46-19.

The following graphic shows the network condition for an SSL rule that blocks encrypted connections originating from your internal network and attempting to access resources either in the Cayman Islands or an offshore holding corporation server at 182.16.0.3.



The example manually specifies the offshore holding corporation's server IP address, and uses a ASA FirePOWER module-provided Cayman Islands geolocation object to represent Cayman Island IP addresses.

You can add a maximum of 50 items to each of the **Source Networks** and **Destination Networks** in a single network condition, and you can mix network and geolocation-based configurations:

- To match encrypted traffic *from* an IP address or geographical location, configure the **Source Networks**.
- To match encrypted traffic *to* an IP address or geographical location, configure the **Destination Networks**.

If you add both source and destination network conditions to a rule, matching encrypted traffic must originate from one of the specified IP addresses **and** be destined for one of the destination IP addresses.

When building a network condition, warning icons indicate invalid configurations. For details, hover your pointer over the icon.

To control traffic by network or geographical location:

Access: Admin/Access Admin/Network Admin

Step 1 In the SSL policy where you want to control encrypted traffic by network, create a new SSL rule or edit an existing rule.

For detailed instructions, see Understanding and Creating SSL Rules, page 16-4.

Step 2 In the SSL rule editor, select the Networks tab.

The Networks tab appears.

Step 3 Find and select the networks you want to add from the **Available Networks**, as follows:

- Click the Networks tab to display network objects and groups to add; click the Geolocation tab to display geolocation objects.

- To add a network object on the fly, which you can then add to the condition, click the add icon ( ) above the Available Networks list; see Working with Network Objects, page 2-3.
- To search for network or geolocation objects to add, select the appropriate tab, click the **Search by name or value** prompt above the **Available Networks** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.

To select an object, click it. To select multiple objects, use the Shift and Ctrl keys, or right-click and then select **Select All**.

Step 4 Click **Add to Source** or **Add to Destination** to add the selected objects to the appropriate list.

You can also drag and drop selected objects.

Step 5 Add any source or destination IP addresses or address blocks that you want to specify manually.

Click the **Enter an IP address** prompt below the **Source Networks** or **Destination Networks** list; then type an IP address or address block and click **Add**.

Step 6 Save or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Controlling Encrypted Traffic by Port

License: Any

Port conditions in SSL rules allow you to control encrypted traffic by its source and destination TCP port. When you build a port-based SSL rule condition, you can manually specify TCP ports. Alternately, you can configure port conditions with port *objects*, which are reusable and associate a name with one or more ports.

Tip After you create a port object, you can use it not only to build SSL rules, but also to represent ports in various other places in the module interface. You can create port objects either using the object manager or on-the-fly while you are configuring SSL rules. For more information, see Working with Port Objects, page 2-9.

You can add a maximum of 50 items to each of the **Selected Source Ports** and **Selected Destination Ports** lists in a single network condition:

- To match encrypted traffic *from* a TCP port, configure the **Selected Source Ports**.
- To match encrypted traffic *to* a TCP port, configure the **Selected Destination Ports**.
- To match encrypted traffic both originating from TCP **Selected Source Ports** *and* destined for TCP **Selected Destination Ports**, configure both.

You can only configure the **Selected Source Ports** and **Selected Destination Ports** lists with TCP ports. Port objects containing non-TCP ports are greyed out in the **Available Ports** list.

When building a port condition, warning icons indicate invalid configurations. For example, you can use the object manager to edit in-use port objects so that the rules that use those object groups become invalid. For details, hover your pointer over the icon.

To control traffic by port:

Step 1 In the SSL policy where you want to control encrypted traffic by TCP port, create a new SSL rule or edit an existing rule.

For detailed instructions, see Understanding and Creating SSL Rules, page 16-4.

Step 2 In the SSL rule editor, select the Ports tab.

The Ports tab appears.

Step 3 Find and select the TCP ports you want to add from the **Available Ports**, as follows:

- To add a TCP port object on the fly, which you can then add to the condition, click the add icon ( ) above the Available Ports list; see Working with Port Objects, page 2-9.
- To search for TCP-based port objects and groups to add, click the **Search by name or value** prompt above the **Available Ports** list, then type either the name of the object, or the value of a port in the object. The list updates as you type to display matching objects. For example, if you type 443, the ASA FirePOWER module displays the ASA FirePOWER module-provided HTTPS port object.

To select a TCP-based port object, click it. To select multiple TCP-based port objects, use the Shift and Ctrl keys, or right-click and then select **Select All**. If the object includes non-TCP-based ports, you cannot add it to your port condition.

Step 4 Click **Add to Source** or **Add to Destination** to add the selected objects to the appropriate list.

You can also drag and drop selected objects.

Step 5 Enter a **Port** under the **Selected Source Ports** or **Selected Destination Ports** list to manually specify source or destination ports. You can specify a single port with a value from 0 to 65535.

Step 6 Click **Add**.

Note that the ASA FirePOWER module will not add a port to a rule condition that results in an invalid configuration.

Step 7 Save or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Controlling Encrypted Traffic Based on User

License: Control

You can configure SSL rules to match traffic for users retrieved from a Microsoft Active Directory Server. User conditions in SSL rules allow you perform *user control*—to manage which traffic can traverse your network, by limiting traffic based on the LDAP user logged into a host.

User control works by associating *access controlled users* with IP addresses. Deployed agents monitor specified users as they log in and out of hosts or authenticate with Active Directory credentials for other reasons. For example, your organization may use services or applications that rely on Active Directory for centralized authentication.

For traffic to match an SSL rule with a user condition, the IP address of either the source or destination host in the monitored session must be associated with a logged in access controlled user. You can control traffic based on individual users or the groups those users belong to.

You can combine user conditions with each other and with other types of conditions to create an SSL rule. These SSL rules can be simple or complex, matching and inspecting traffic using multiple conditions. For detailed information on SSL rules, see Understanding and Creating SSL Rules, page 16-4.

User control requires a Control license and is supported only for LDAP users and groups (*access controlled users*), using login and logoff records reported by a User Agent monitoring Microsoft Active Directory servers.

Before you can write SSL rules with user conditions, you must configure a connection between the ASA FirePOWER module and at least one of your organization's Microsoft Active Directory servers. This configuration, called an authentication object, contains connection settings and authentication filter settings for the server. It also specifies the users you can use in user conditions.

In addition, you must install User Agents. The agents monitor users when they authenticate against Active Directory credentials, and send records of those logins to the ASA FirePOWER module. These records associate users with IP addresses, which is what allows SSL rules with user conditions to trigger.

To control encrypted traffic by user:

Step 1 In the SSL policy where you want to control encrypted traffic by user, create a new SSL rule or edit an existing rule.

For detailed instructions, see Understanding and Creating SSL Rules, page 16-4.

Step 2 In the SSL rule editor, select the Users tab.

The Users tab appears.

Step 3 To search for users to add, click the **Search by name or value** prompt above the **Available Users** list, then type the username. The list updates as you type to display matching users.

To select a user, click it. To select multiple users, use the Shift and Ctrl keys, or right-click and then select **Select All**.

Step 4 Click **Add to Rule** or to add the selected users to the **Selected Users** list.

You can also drag and drop selected users.

Step 5 Save or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Controlling Encrypted Traffic by Reputation

License: Control or URL Filtering

Reputation-based conditions in SSL rules allow you to manage which encrypted traffic can traverse your network, by contextualizing your network traffic and limiting it where appropriate. SSL rules govern the following types of reputation-based control:

- Application conditions allow you to perform *application control*, which controls application traffic based on not only individual applications, but also applications' basic characteristics: type, risk, business relevance, and categories.
- URL conditions allow you to control web traffic based on a websites' assigned category and reputation.

You can combine reputation-based conditions with each other and with other types of conditions to create an SSL rule. These SSL rules can be simple or complex, matching and inspecting traffic using multiple conditions.

For more information, see the following sections:

- Controlling Encrypted Traffic Based on Application, page 17-8
- Controlling Encrypted Traffic by URL Category and Reputation, page 17-13

Controlling Encrypted Traffic Based on Application

License: Control

When the Firepower system analyzes encrypted IP traffic, it can identify and classify commonly used encrypted applications on your network prior to decrypting the encrypted session. The ASA FirePOWER module uses this discovery-based *application awareness* feature to allow you to control encrypted application traffic on your network.

Application conditions in SSL rules allow you to perform this *application control*. Within a single SSL rule, there are a few ways you can specify applications whose traffic you want to control:

- You can select individual applications, including custom applications.
- You can use ASA FirePOWER module-provided *application filters*, which are named sets of applications organized according to its basic characteristics: type, risk, business relevance, and categories.
- You can create and use custom application filters, which group applications (including custom applications) in any way you choose.

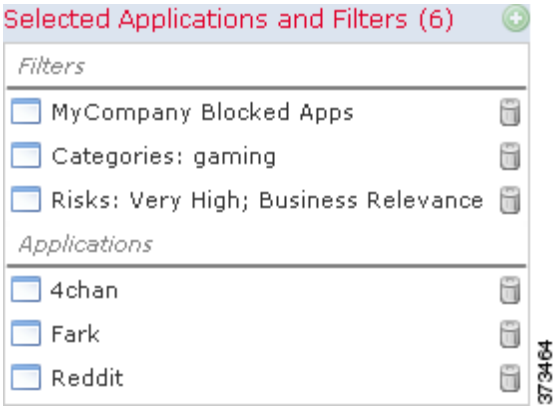
Note When you filter application traffic using access control rules, you can use application tags as a criterion. to filter. However, you cannot use application tags to filter encrypted traffic because there is no benefit. All applications that the ASA FirePOWER module can detect in encrypted traffic are tagged **SSL Protocol**; applications without this tag can only be detected in unencrypted or decrypted traffic.

Application filters allow you to quickly create application conditions for SSL rules. They simplify policy creation and administration, and grant you assurance that the module will control web traffic as expected. For example, you could create an SSL rule that identifies and decrypts all high risk, low business relevance applications in encrypted traffic. If a user attempts to use one of those applications, the session is decrypted and inspected with access control.

In addition, Cisco frequently updates and adds additional detectors via system and vulnerability database (VDB) updates. You can also create your own detectors and assign characteristics (risk, relevance, and so on) to the applications they detect. By using filters based on application characteristics, you can ensure that the module uses the most up-to-date detectors to monitor application traffic.

For traffic to match an SSL rule with an application condition, the traffic must match one of the filters or applications that you add to a **Selected Applications and Filters** list.

The following graphic shows the application condition for an SSL rule that decrypts a custom group of applications for MyCompany, all applications with high risk and low business relevance, gaming applications, and some individually selected applications.



In a single application condition, you can add a maximum of 50 items to the **Selected Applications and Filters** list. Each of the following counts as an item:

- One or more filters from the **Application Filters** list, individually or in custom combination. This item represents set of applications, grouped by characteristic.
- A filter created by saving search of the applications in the **Available Applications** list. This item represents a set of applications, grouped by substring match.
- An individual application from the **Available Applications** list.

In the module interface, filters added to a condition are listed above and separately from individually added applications.

Note that when you apply an SSL policy, for each rule with an application condition, the ASA FirePOWER module generates a list of unique applications to match. In other words, you may use overlapping filters and individually specified applications to ensure complete coverage.

For more information, see the following sections:

- Matching Encrypted Traffic with Application Filters, page 17-9
- Matching Traffic from Individual Applications, page 17-10
- Adding an Application Condition to an SSL Rule, page 17-11
- Limitations to Encrypted Application Control, page 17-12

Matching Encrypted Traffic with Application Filters

License: Control

When building an application condition in an SSL rule, use the **Application Filters** list to create a set of applications, grouped by characteristic, whose traffic you want to match.

For your convenience, the ASA FirePOWER module characterizes each application that it detects using a specified criteria. You can use these criteria as filters or create custom combinations of filters to perform application control.

Note that the mechanism for filtering applications within an SSL rule is the same as that for creating reusable, custom application filters using the object manager; see Working with Application Filters, page 2-11. You can also save many filters you create on-the-fly in access control rules as new, reusable filters. You cannot save a filter that includes another user-created filter because you cannot nest user-created filters.

Understanding How Filters Are Combined

When you select filters, singly or in combination, the **Available Applications** list updates to display only the applications that meet your criteria. You can select ASA FirePOWER module-provided filters in combination, but not custom filters.

The module links multiple filters of the same filter type with an OR operation. For example, if you select the Medium and High filters under the Risks type, the resulting filter is:

```
Risk: Medium OR High
```

If the Medium filter contained 110 applications and the High filter contained 82 applications, the module displays all 192 applications in the **Available Applications** list.

The module links different types of filters with an AND operation. For example, if you select the Medium and High filters under the Risks type, and the Medium and High filters under the Business Relevance type, the resulting filter is:

```
Risk: Medium OR High
AND
Business Relevance: Medium OR High
```

In this case, the module displays only those applications that are included in both the Medium or High Risk type AND the Medium or High Business Relevance type.

Finding and Selecting Filters

To select filters, click the arrow next to a filter type to expand it, then select or clear the check box next to each filter whose applications you want to display or hide. You can also right-click a Cisco-provided filter type (**Risks**, **Business Relevance**, **Types**, or **Categories**) and select **Check All** or **Uncheck All**.

To search for filters, click the **Search by name** prompt above the **Available Filters** list, then type a name. The list updates as you type to display matching filters.

After you are done selecting filters, use the **Available Applications** list to add those filters to the rule; see Matching Traffic from Individual Applications, page 17-10.

Matching Traffic from Individual Applications

License: Control

When building an application condition in an SSL rule, use the **Available Applications** list to select the applications whose traffic you want to match.

Browsing the List of Applications

When you first start to build the condition the list is unconstrained, and displays every application the module detects, 100 at a time:

- To page through the applications, click the arrows underneath the list.
- To display a pop-up window with summary information about the application's characteristics, as well as Internet search links that you can follow, click the information icon ( ) next to an application.

Finding Applications to Match

To help you find the applications you want to match, you can constrain the **Available Applications** list in the following ways:

- To search for applications, click the **Search by name** prompt above the list, then type a name. The list updates as you type to display matching applications.

- To constrain the applications by applying a filter, use the **Application Filters** list (see Matching Encrypted Traffic with Application Filters, page 17-9). The **Available Applications** list updates as you apply filters.

Once constrained, an **All apps matching the filter** option appears at the top of the **Available Applications** list. This option allows you to add all the applications in the constrained list to the **Selected Applications and Filters** list, all at once.

Note If you select one or more filters in the Application Filters list and also search the **Available Applications** list, your selections and the search-filtered **Available Applications** list are combined using an AND operation. That is, the **All apps matching the filter** condition includes all the individual conditions currently displayed in the **Available Applications** list as well as the search string entered above the **Available Applications** list.

Selecting Single Applications to Match in a Condition

After you find an application you want to match, click to select it. To select multiple applications, use the Shift and Ctrl keys, or right-click and select **Select All** to select all applications in the current constrained view.

In a single application condition, you can match a maximum of 50 applications by selecting them individually; to add more than 50 you must either create multiple SSL rules or use filters to group applications.

Selecting All Applications Matching a Filter for a Condition

Once constrained by either searching or using the filters in the **Application Filters** list, the **All apps matching the filter** option appears at the top of the **Available Applications** list.

This option allows you to add the entire set of applications in the constrained **Available Applications** list to the **Selected Applications and Filters** list, at once. In contrast to adding applications individually, adding this set of applications counts as only one item against the maximum of 50, regardless of the number of individual application that comprise it.

When you build an application condition this way, the name of the filter you add to the **Selected Applications and Filters** list is a concatenation of the filter types represented in the filter plus the names of up to three filters for each type. More than three filters of the same type are followed by an ellipsis (...). For example, the following filter name includes two filters under the Risks type and four under Business Relevance:

```
Risks: Medium, High Business Relevance: Low, Medium, High,...
```

Filter types that are not represented in a filter you add with **All apps matching the filter** are not included in the name of the filter you add. The instructional text that is displayed when you hover your pointer over the filter name in the **Selected Applications and Filters** list indicates that these filter types are set to *any*; that is, these filter types do not constrain the filter, so any value is allowed for these.

You can add multiple instances of **All apps matching the filter** to an application condition, with each instance counting as a separate item in the **Selected Applications and Filters** list. For example, you could add all high risk applications as one item, clear your selections, then add all low business relevance applications as another item. This application condition matches applications that are high risk OR have low business relevance.

Adding an Application Condition to an SSL Rule

License: Control

For encrypted traffic to match an SSL rule with an application condition, the traffic must match one of the filters or applications that you add to a **Selected Applications and Filters** list.

You can add a maximum of 50 items per condition, and filters added to a condition are listed above and separately from individually added applications. When building an application condition, warning icons indicate invalid configurations. For details, hover your pointer over the icon.

To control encrypted application traffic:

Step 1 In the SSL policy where you want to control traffic by application, create a new SSL rule or edit an existing rule.

For detailed instructions, see Understanding and Creating SSL Rules, page 16-4.

Step 2 In the SSL rule editor, select the Applications tab.

The Applications tab appears.

Step 3 Optionally, use filters to constrain the list of applications displayed in the **Available Applications** list.

Select one or more filters in the **Application Filters** list. For more information, see Matching Encrypted Traffic with Application Filters, page 17-9.

Step 4 Find and select the applications you want to add from the **Available Applications** list.

You can search for and select individual applications, or, when the list is constrained, **All apps matching the filter**. For more information, see Matching Traffic from Individual Applications, page 17-10.

Step 5 Click **Add to Rule** to add the selected applications to the **Selected Applications and Filters** list.

You can also drag and drop selected applications and filters. Filters appear under the heading *Filters*, and applications appear under the heading *Applications*.



Tip Before you add another filter to this application condition, click **Clear All Filters** to clear your existing selections.

Step 6 Save or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Limitations to Encrypted Application Control

License: Control

Keep the following points in mind when performing application control.

Encrypted Application Identification

The ASA FirePOWER module can identify unencrypted applications that become encrypted using StartTLS. This includes such applications as SMTPS, POPS, FTPS, TelnetS, and IMAPS. In addition, it can identify certain encrypted applications based on the Server Name Indication in the TLS client hello message, or the server certificate subject distinguished name value.

Speed of Application Identification

The ASA FirePOWER module cannot perform application control on encrypted traffic before:

- an encrypted connection is established between a client and server, and
- the module identifies the application in the encrypted session

This identification occurs after the server certificate exchange. If traffic exchanged during the handshake matches all other conditions in an SSL rule containing an application condition but the identification is not complete, the SSL policy allows the packet to pass. This behavior allows the handshake to complete so that applications can be identified. For your convenience, affected rules are marked with an information icon ( ).

After the module completes its identification, it applies the SSL rule action to the remaining session traffic that matches its application condition.

Controlling Encrypted Traffic by URL Category and Reputation

License: URL Filtering

URL conditions in SSL rules allow you to handle and decrypt encrypted website traffic that users on your network can access. The module detects the requested URL based on information passed during the SSL handshake. With a URL Filtering license, you can control access to websites based on the URL's general classification, or *category*, and risk level, or *reputation*.

Note You can handle and decrypt traffic to specific URLs by defining a distinguished name SSL rule condition. The common name attribute in a certificate's subject distinguished name contains the site's URL. For more information, see Controlling Encrypted Traffic by Certificate Distinguished Name, page 17-17.

For more information, see:

- Performing Reputation-Based URL Blocking, page 17-13
- Limitations on URL Detection and Blocking, page 17-16

Performing Reputation-Based URL Blocking

License: URL Filtering

With a URL Filtering license, you can control your users' access to websites based on the category and reputation of requested URLs:

- The URL *category* is a general classification for the URL. For example, ebay.com belongs to the **Auctions** category, and monster.com belongs to the **Job Search** category. A URL can belong to more than one category.
- The URL *reputation* represents how likely the URL is to be used for purposes that might be against your organization's security policy. A URL's risk can range from **High Risk** (level 1) to **Well Known** (level 5).

URL categories and reputations, which the Firepower system obtains from the Cisco cloud, allow you to quickly create URL conditions for SSL rules. For example, you could create an SSL rule that identifies and blocks all **High risk** URLs in the **Abused Drugs** category. If a user attempts to browse to any URL with that category and reputation combination over an encrypted connection, the session is blocked.

Note Before SSL rules with category and reputation-based URL conditions can take effect, you **must** enable communications with the Cisco cloud. This allows the ASA FirePOWER module to retrieve URL data. For more information, see Enabling Cloud Communications, page 44-2.

Using category and reputation data from the Cisco cloud simplifies policy creation and administration. It grants you assurance that the module controls encrypted web traffic as expected. Finally, because the cloud is continually updated with new URLs, as well as new categories and risks for existing URLs, you can ensure that the module uses up-to-date information to filter requested URLs. Malicious sites that represent security threats such as malware, spam, botnets, and phishing may appear and disappear faster than you can update and apply new policies.

For example:

- If a rule blocks all gaming sites, as new domains get registered and classified as **Gaming**, the module can block those sites automatically.
- If a rule blocks all malware, and a blog page gets infected with malware, the cloud can recategorize the URL from **Blog** to **Malware** and the module can block that site.
- If a rule blocks high-risk social networking sites, and somebody posts a link on their profile page that contains links to malicious payloads, the cloud can change the reputation of that page from **Benign sites** to **High risk** so the module can block it.

Note that if the cloud does not know the category or reputation of a URL, or if the ASA FirePOWER module cannot contact the cloud, that URL does **not** trigger SSL rules with category or reputation-based URL conditions. You cannot assign categories or reputations to URLs manually.

The following graphic shows the URL condition for an access control rule that blocks: all malware sites, all high-risk sites, and all non-benign social networking sites.



Tip If you decrypt traffic, then block it with access control, you can give users a chance to bypass the block by clicking through a warning page. See Interactive Blocking Actions: Allowing Users to Bypass Website Blocks, page 6-9 for more information.

You can add a maximum of 50 **Selected Categories** to match in a single URL condition. Each URL category, optionally qualified by reputation, counts as a single item.

The following table summarizes how you build the condition shown above. Note that you cannot qualify a literal URL or URL object with a reputation.

Table 17-1 Example: Building A URL Condition

To block...	Select this Category or URL Object...	And this Reputation...
malware sites, regardless of reputation	Malware Sites	Any
any URL with a high risk (level 1)	Any	1 - High Risk
social networking sites with a risk greater than benign (levels 1 through 3)	Social Network	3 - Benign sites with security risks

When building a URL condition, warning icons indicate invalid configurations. For details, hover your pointer over the icon and see Troubleshooting Access Control Policies and Rules, page 4-12.

To control traffic by requested URL using category and reputation data:

Step 1 In the SSL policy where you want to control encrypted traffic by URL, create a new SSL rule or edit an existing rule.

For detailed instructions, see Understanding and Creating SSL Rules, page 16-4.

Step 2 In the SSL rule editor, select the Categories tab.

The Categories tab appears.

Step 3 Find and select the categories of URL you want to add from the **Categories** list. To match encrypted web traffic regardless of category, select **Any** category.

To search for categories to add, click the **Search by name or value** prompt above the **Categories** list, then type the category name. The list updates as you type to display matching categories.

To select a category, click it. To select multiple categories, use the Shift and Ctrl keys.

Tip Although you can right-click and **Select All** categories, adding all categories this way exceeds the 50-item maximum for an SSL rule. Instead, use **Any**.

Step 4 Optionally, qualify your category selections by clicking a reputation level from the **Reputations** list. If you do not specify a reputation level, the module defaults to **Any**, meaning all levels.

You can only select one reputation level. When you choose a reputation level, the SSL rule behaves differently depending on its purpose:

- If the rule blocks web access or decrypts traffic (the rule action is **Block**, **Block with reset**, **Decrypt - Known Key**, **Decrypt - Resign**, or **Monitor**) selecting a reputation level also selects all reputations more severe than that level. For example, if you configure a rule to block **Suspicious sites** (level 2), it also automatically blocks **High Risk** (level 1) sites.
- If the rule allows web access, subject to access control (the rule action is **Do not decrypt**), selecting a reputation level also selects all reputations less severe than that level. For example, if you configure a rule to allow **Benign sites** (level 4), it also automatically allows **Well known** (level 5) sites.

If you change the rule action for a rule, the module automatically changes the reputation levels in URL conditions according to the above points.

Step 5 Click **Add to Rule** or to add the selected items to the **Selected Categories** list.

You can also drag and drop selected items.

Step 6 Save or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Limitations on URL Detection and Blocking

License: URL Filtering

Keep the following points in mind when performing URL detection and blocking.

Speed of URL Identification

The module cannot categorize URLs before:

- a monitored connection is established between a client and server
- the module identifies the HTTPS application in the session
- the module identifies the requested URL from either the client hello message or the server certificate

This identification occurs after the server certificate exchange. If traffic exchanged during the handshake matches all other conditions in an SSL rule containing a URL condition but the identification is not complete, the SSL policy allows the packet to pass. This behavior allows the connection to be established so that URLs can be identified. For your convenience, affected rules are marked with an information icon ( ).

After the module completes its identification, it applies the SSL rule action to the remaining session traffic that matches its URL condition.

Search Query Parameters in URLs

The module does not use search query parameters in the URL to match URL conditions. For example, consider a scenario where you block all shopping traffic. In that case, using a web search to search for amazon.com is not blocked, but browsing to amazon.com is.

Controlling Traffic Based on Server Certificate Characteristics

License: Any

You can create SSL rules that handle and decrypt encrypted traffic based on server certificate characteristics. You can detect the protocol version or cipher suite used to encrypt the session, and handle traffic accordingly. You can also detect the server certificate and handle traffic, based on the following server certificate characteristics:

- the server certificate itself
- the certificate issuer, whether an issuing CA or if the certificate is self-signed
- the certificate holder
- various certificate statuses, such as whether the certificate is valid, or revoked by the issuing CA

To detect multiple cipher suites in a rule, the certificate issuer, or the certificate holder, you can create reusable cipher suite list and distinguished name objects and add them to your rule. To detect the server certificate and certain certificate statuses, you must create external certificate and external CA objects for the rule.

For more information, see the following sections:

- Controlling Encrypted Traffic by Certificate Distinguished Name, page 17-17
- Controlling Encrypted Traffic by Certificate, page 17-19
- Controlling Encrypted Traffic by Certificate Status, page 17-20
- Controlling Encrypted Traffic by Cipher Suite, page 17-25
- Controlling Traffic by Encryption Protocol Version, page 17-26

Controlling Encrypted Traffic by Certificate Distinguished Name

License: Any

Distinguished name conditions in SSL rules allow you to handle and inspect encrypted traffic based on the CA that issued a server certificate, or the certificate holder. Based on the issuer distinguished name, you can handle traffic based on the CA that issued a site's server certificate.

When configuring the rule condition, you can manually specify a literal value, reference a distinguished name object, or reference a distinguished name group containing multiple objects.

Note You cannot configure a distinguished name condition if you also select the **Decrypt - Known Key** action. Because that action requires you to select a server certificate to decrypt traffic, the certificate already matches the traffic. See Decrypt Actions: Decrypting Traffic for Further Inspection, page 16-10 for more information.

You can match against multiple subject and issuer distinguished names in a single certificate status rule condition; only one common or distinguished name needs to match to match the rule.

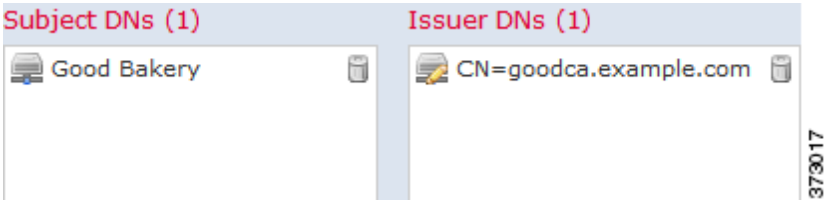
If you add a distinguished name manually, it can contain the common name attribute (**CN**). If you add a common name without `CN=` then the module prepends `CN=` before saving the object.

You can also add a distinguished name with one of each attribute listed in the following table, separated by commas.

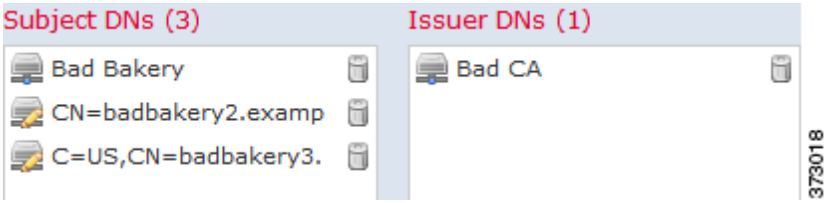
Table 17-2 Distinguished Name Attributes

Attribute	Description	Allowed Values
C	Country Code	two alphabetic characters
CN	Common Name	up to 64 alphanumeric, backslash (\), hyphen (-), quotation ("), asterisk (*), period (.), or space characters
O	Organization	
OU	Organizational Unit	

The following graphic illustrates a distinguished name rule condition searching for certificates issued to goodbakery.example.com or issued by goodca.example.com. Traffic encrypted with these certificates is allowed, subject to access control.



The following graphic illustrates a distinguished name rule condition searching for certificates issued to badbakery.example.com and associated domains, or certificates issued by badca.example.com. Traffic encrypted with these certificates is decrypted using a re-signed certificate.



You can add a maximum of 50 literal values and distinguished name objects to the **Subject DNs**, and 50 literal values and distinguished name objects to the **Issuer DNs**, in a single DN condition.

The ASA FirePOWER module-provided DN object group, Sourcefire Undecryptable Sites, contains websites whose traffic the module cannot decrypt. You can add this group to a DN condition to block or not decrypt traffic to or from these websites, without wasting system resources attempting to decrypt that traffic. You can modify individual entries in the group. You cannot delete the group. System updates can modify the entries on this list, but the module preserves user changes.

The first time the system detects an encrypted session to a new server, DN data is not available for ClientHello processing, which can result in an undecrypted first session. After the initial session, the managed device caches data from the server Certificate message. For subsequent connections from the same client, the system can match the ClientHello message conclusively to rules with DN conditions and process the message to maximize decryption potential.

To inspect encrypted traffic based on certificate subject or issuer distinguished name:

Step 1 In the SSL policy where you want to control encrypted traffic by certificate subject or issuer distinguished name, create a new SSL rule or edit an existing rule.

For detailed instructions, see Understanding and Creating SSL Rules, page 16-4.

Step 2 In the SSL rule editor, select the DN tab.

The DN tab appears.

Step 3 Find and select the distinguished names you want to add from the **Available DNs**, as follows:

- To add a distinguished name object on the fly, which you can then add to the condition, click the add icon ( ) above the **Available DNs** list; see Working with Distinguished Name Objects, page 2-33.

- To search for distinguished name objects and groups to add, click the **Search by name or value** prompt above the **Available DNs** list, then type either the name of the object, or a value in the object. The list updates as you type to display matching objects.

To select an object, click it. To select multiple objects, use the Shift and Ctrl keys, or right-click and then select **Select All**.

Step 4 You have the following options:

- Click **Add to Subject** to add the selected objects to the **Subject DNs** list.
- Click **Add to Issuer** to add the selected objects to the **Issuer DNs** list.

You can also drag and drop selected objects.

Step 5 Add any literal common names or distinguished names that you want to specify manually.

Click the **Enter DN or CN** prompt below the **Subject DNs** or **Issuer DNs** list; then type a common name or distinguished name and click **Add**.

Step 6 Add or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Controlling Encrypted Traffic by Certificate

License: Any

Certificate conditions in SSL rules allow you to handle and inspect encrypted traffic based on the server certificate used to encrypt that traffic. You can configure a condition with one or more certificates; traffic matches the rule if the certificate matches any of the condition's certificates.

When you build a certificate-based SSL rule condition, you can upload a server certificate; you save the certificate as an external certificate *object*, which is reusable and associates a name with a server certificate. Alternately, you can configure certificate conditions with existing external certificate objects and object groups.

You can search the **Available Certificates** field in the rule condition based for external certificate objects and object groups based on the following certificate distinguished name characteristics:

- subject or issuer common name (CN)
- subject or issuer organization (O)
- subject or issuer organizational unit (OU)

You can choose to match against multiple certificates in a single certificate rule condition; if the certificate used to encrypt the traffic matches any of the uploaded certificates, the encrypted traffic matches the rule.

You can add a maximum of 50 external certificate objects and external certificate object groups to the **Selected Certificates** in a single certificate condition.

Note the following:

- You cannot configure a certificate condition if you also select the **Decrypt - Known Key** action. Because that action requires you to select a server certificate to decrypt traffic, the implication is that the certificate already matches the traffic. See Decrypt Actions: Decrypting Traffic for Further Inspection, page 16-10 for more information.

- If you configure a certificate condition with an external certificate object, any cipher suites you add to a cipher suite condition, or internal CA objects you associate with the **Decrypt - Resign** action, must match the external certificate's signature algorithm type. For example, if your rule's certificate condition references an EC-based server certificate, any cipher suites you add, or CA certificates you associate with the **Decrypt - Resign** action, must also be EC-based. If you mismatch signature algorithm types in this case, the policy editor displays a warning icon next to the rule. For more information, see Controlling Encrypted Traffic by Cipher Suite, page 17-25 and Decrypt Actions: Decrypting Traffic for Further Inspection, page 16-10.
- The first time the system detects an encrypted session to a new server, certificate data is not available for ClientHello processing, which can result in an undecrypted first session. After the initial session, the managed device caches data from the server Certificate message. For subsequent connections from the same client, the system can match the ClientHello message conclusively to rules with certificate conditions and process the message to maximize decryption potential.

To inspect encrypted traffic based on server certificate:

Step 1 In the SSL policy where you want to control encrypted traffic based on server certificate, create a new SSL rule or edit an existing rule.

For detailed instructions, see Understanding and Creating SSL Rules, page 16-4.

Step 2 In the SSL rule editor, select the Certificate tab.

The Certificate tab appears.

Step 3 Find and select the server certificates you want to add from the **Available Certificates**, as follows;

- To add an external certificate object on the fly, which you can then add to the condition, click the add icon ( ) above the **Available Certificates** list; see Working with External Certificate Objects, page 2-41.
- To search for certificate objects and groups to add, click the **Search by name or value** prompt above the **Available Certificates** list, then type either the name of the object, or a value in the object. The list updates as you type to display matching objects.

To select an object, click it. To select multiple objects, use the Shift and Ctrl keys, or right-click and then select **Select All**.

Step 4 Click **Add to Rule** to add the selected objects to the **Subject Certificates** list.

You can also drag and drop selected objects.

Step 5 Add or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Controlling Encrypted Traffic by Certificate Status

License: Any

Certificate status conditions in SSL rules allow you to handle and inspect encrypted traffic based on the status of the server certificate used to encrypt the traffic, including whether a certificate is valid, revoked, expired, not yet valid, self-signed, or signed by a trusted CA.

Checking whether a CA issued or revoked a certificate requires uploading root and intermediate CA certificates and associated CRLs as objects. You then add these trusted CA objects to an SSL policy's list of trusted CA certificates.

For each certificate status SSL rule condition you configure, you can match traffic against the presence or absence of a given status. You can select several statuses in one rule condition; if the certificate matches any of the selected statuses, the rule matches the traffic.

For more information, see:

- Trusting External Certificate Authorities, page 17-21
- Matching Traffic on Certificate Status, page 17-22

Trusting External Certificate Authorities

License: Any

You can trust CAs by adding root and intermediate CA certificates to your SSL policy, then use these trusted CAs to verify server certificates used to encrypt traffic. Verified server certificates include certificates signed by trusted CAs.

If a trusted CA certificate contains an uploaded certificate revocation list (CRL), you can also verify whether a trusted CA revoked the encryption certificate. See Adding a Certificate Revocation List to a Trusted CA Object, page 2-40 for more information.

After you add trusted CA certificates to the SSL policy, you can configure an SSL rule with various Certificate Status conditions to match against this traffic. See Working with Trusted Certificate Authority Objects, page 2-39 and Controlling Encrypted Traffic by Certificate Status, page 17-20 for more information.

Tip Upload all certificates within a root CA's chain of trust to the list of trusted CA certificates, including the root CA certificate and all intermediate CA certificates. Otherwise, it is more difficult to detect trusted certificates issued by intermediate CAs.

When you create an SSL policy, the ASA FirePOWER module populates the Trusted CA Certificates tab with a default Trusted CA object group, Cisco Trusted Authorities. You can modify individual entries in the group, and choose whether to include this group in your SSL policy. You cannot delete the group. System updates can modify the entries on this list, but user changes are preserved. See Creating a Basic SSL Policy, page 15-2 for more information.

To add trusted CAs to your policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > SSL**.

The SSL Policy page appears.

Step 2 Click the edit icon ( ) next to the SSL policy you want to configure.

The SSL policy editor appears.

Step 3 Select the **Trusted CA Certificates** tab.

The Trusted CA Certificates page appears.

Step 4 Find and select the trusted CAs you want to add from the **Available Trusted CAs**, as follows:

- To add a trusted CA object on the fly, which you can then add to the condition, click the add icon ( ) above the **Available Trusted CAs** list; see Working with Trusted Certificate Authority Objects, page 2-39.
- To search for trusted CA objects and groups to add, click the **Search by name or value** prompt above the **Available Trusted CAs** list, then type either the name of the object, or a value in the object. The list updates as you type to display matching objects.

To select an object, click it. To select multiple objects, use the Shift and Ctrl keys, or right-click and then select **Select All**.

Step 5 Click **Add to Rule** to add the selected objects to the **Selected Trusted CAs** list.

You can also drag and drop selected objects.

Step 6 Add or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Matching Traffic on Certificate Status

License: Any

Based on the certificate status rule condition configuration, you can match encrypted traffic based on the status of the server certificate used to encrypt traffic. You can:

- check for a server certificate status
- check that a certificate does not have a status
- skip checking for the presence or absence of a certificate status

You can choose to match against the presence or absence of multiple certificate statuses in a single certificate status rule condition; the certificate needs to only match one of the criteria to match the rule.

The following table describes how the ASA FirePOWER module evaluates encrypted traffic based on the encrypting server certificate's status.

Table 17-3 Certificate Status Rule Condition Criteria

Status Check	Status Set to Yes	Status Set to No
Revoked	The policy trusts the CA that issued the server certificate, and the CA certificate uploaded to the policy contains a CRL that revokes the server certificate.	The policy trusts the CA that issued the server certificate, and the CA certificate uploaded to the policy does not contain a CRL that revokes the certificate.
Self-signed	The detected server certificate contains the same subject and issuer distinguished name.	The detected server certificate contains different subject and issuer distinguished names.

Table 17-3 Certificate Status Rule Condition Criteria (continued)

Status Check	Status Set to Yes	Status Set to No
Valid	All of the following are true: • The policy trusts the CA that issued the certificate • The signature is valid • The issuer is valid • None of the policy's trusted CAs revoked the certificate. • The current date is between the certificate Valid From and Valid To date	At least one of the following is true: • The policy does not trust the CA that issued the certificate • The signature is invalid • The issuer is invalid • A trusted CA in the policy revoked the certificate • The current date is before the certificate Valid From date • The current date is after the certificate Valid To date
Invalid signature	The certificate's signature cannot be properly validated against the certificate's content.	The certificate's signature is properly validated against the certificate's content.
Invalid issuer	The issuer CA certificate is not stored in the policy's list of trusted CA certificates.	The issuer CA certificate is stored in the policy's list of trusted CA certificates.
Expired	The current date is after the certificate Valid To date.	The current date is before or on the certificate Valid To date.
Not yet valid	The current date is before the certificate Valid From date.	The current date is after or on the certificate Valid From date.

Consider the following example. The organization trusts the Verified Authority certificate authority. The organization does not trust the Spammer Authority certificate authority. The system administrator uploads the Verified Authority certificate and an intermediate CA certificate issued by Verified Authority to the module. Because Verified Authority revoked a certificate it previously issued, the system administrator uploads the CRL that Verified Authority distributed.

The following graphic illustrates a certificate status rule condition checking for valid certificates, those issued by Verified Authority, not on the CRL, and still within the Valid From and Valid To date. Because of the configuration, traffic encrypted with these certificates is not decrypted and inspected with access control.



The following graphic illustrates a certificate status rule condition checking for the absence of a status. In this case, because of the configuration, it matches against traffic encrypted with a certificate that has not expired and monitors that traffic.



The following graphic illustrates a certificate status rule condition that matches on the presence or absence of several statuses. Because of the configuration, if the rule matches incoming traffic encrypted with a certificate issued by an invalid user, self-signed, invalid, or expired, it decrypts the traffic with a known key.



Note that even though a certificate may match more than one status, the rule only takes an action on the traffic once.

Note The first time the system detects an encrypted session to a new server, certificate status is not available for ClientHello processing, which can result in an undecrypted first session. After the initial session, the managed device caches data from the server Certificate message. For subsequent connections from the same client, the system can match the ClientHello message conclusively to rules with certificate status conditions and process the message to maximize decryption potential.

To inspect encrypted traffic by server certificate status:

Step 1 In the SSL policy where you want to control encrypted traffic based on server certificate status, create a new SSL rule or edit an existing rule.

For detailed instructions, see Understanding and Creating SSL Rules, page 16-4.

Step 2 In the SSL rule editor, select the Cert Status tab.

The Cert Status tab appears.

Step 3 For each certificate status, you have the following options:

- Select **Yes** to match against the presence of that certificate status.
- Select **No** to match against the absence of that certificate status.
- Select **Do Not Match** to not match that certificate status.

Step 4 Add or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Controlling Encrypted Traffic by Cipher Suite

License: Any

Cipher suite conditions in SSL rules allow you to handle and inspect encrypted traffic based on the cipher suite used to negotiate the encrypted session. Cisco provides predefined cipher suites you can add to a cipher suite rule condition. You can also add cipher suite list objects containing multiple cipher suites. For more information on cipher suite lists, see Working with Geolocation Objects, page 2-42.

Note You cannot add new cipher suites. You can neither modify nor delete predefined cipher suites.

You can add a maximum of 50 cipher suites and cipher suite lists to the **Selected Cipher Suites** in a single Cipher Suite condition.

Note the following:

- If you add cipher suites not supported for your deployment, you cannot apply the access control policy associated with the SSL policy. For example, passive deployments do not support decrypting traffic with the any of the ephemeral Diffie-Hellman (DHE) or ephemeral elliptic curve Diffie-Hellman (ECDHE) cipher suites. Creating a rule with these cipher suites prevents you from applying your access control policy.
- If you configure a cipher suite condition with a cipher suite, any external certificate objects you add to a certificate condition, or internal CA objects you associate with the **Decrypt - Resign** action, must match the cipher suite's signature algorithm type. For example, if your rule's cipher suite condition references an EC-based cipher suite, any server certificates you add, or CA certificates you associate with the **Decrypt - Resign** action, must also be EC-based. If you mismatch signature algorithm types in this case, the policy editor displays a warning icon next to the rule. For more information, see Controlling Encrypted Traffic by Cipher Suite, page 17-25 and Decrypt Actions: Decrypting Traffic for Further Inspection, page 16-10.
- You can add an anonymous cipher suite to the Cipher Suite condition in an SSL rule, but keep in mind:
  - The system automatically strips anonymous cipher suites during ClientHello processing. For the system to use the rule, you must also configure your SSL rules in an order that prevents ClientHello processing. For more information, see Ordering SSL Rules to Improve Performance and Avoid Preemption, page 16-16.
  - You cannot use either the **Decrypt - Resign** or **Decrypt - Known Key** action in the rule, because the system cannot decrypt traffic encrypted with an anonymous cipher suite.

- When specifying a cipher suite as a rule condition, consider that the rule matches on the negotiated cipher suite in the ServerHello message, rather than on the full list of cipher suites specified in the ClientHello message. During ClientHello processing, the managed device strips unsupported cipher suites from the ClientHello message. However, if this results in all specified cipher suites being stripped, the system retains the original list. If the system retains unsupported cipher suites, subsequent evaluation results in an undecrypted session.

To inspect encrypted traffic by cipher suite:

Step 1 In the SSL policy where you want to control encrypted traffic by cipher suite, create a new SSL rule or edit an existing rule.

For detailed instructions, see Understanding and Creating SSL Rules, page 16-4.

Step 2 In the SSL rule editor, select the Cipher Suite tab.

The Cipher Suite tab appears.

Step 3 Find and select the cipher suites you want to add from the **Available Cipher Suites**, as follows;

- To add a cipher suite list on the fly, which you can then add to the condition, click the add icon ( ) above the **Available Cipher Suites** list; see Working with Geolocation Objects, page 2-42.
- To search for cipher suites and lists to add, click the **Search by name or value** prompt above the **Available Cipher Suites** list, then type either the name of the cipher suite, or a value in the cipher suite. The list updates as you type to display matching cipher suites.

To select a cipher suite, click it. To select multiple cipher suites, use the Shift and Ctrl keys, or right-click and then select **Select All**.

Step 4 Click **Add to Rule** to add the selected cipher suites to the **Selected Cipher Suites** list.

You can also drag and drop selected cipher suites.

Step 5 Add or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Controlling Traffic by Encryption Protocol Version

License: Any

Session conditions in SSL rules allow you to inspect encrypted traffic based on the SSL or TLS version used to encrypt the traffic. You can choose to match against traffic encrypted with SSL version 3.0, or TLS version 1.0, 1.1, or 1.2. By default, all protocol versions are selected when you create a rule; if you select multiple versions, encrypted traffic that matches any of the selected versions matches the rule. You must select at least one protocol version when saving the rule condition.

Note You cannot select SSL v2.0 in a version rule condition; the ASA FirePOWER module does not support decrypting traffic encrypted with SSL version 2.0. You can configure an undecryptable action to allow or block this traffic without further inspection. For more information, see Logging Decryptable Connections with SSL Rules, page 36-14.

To inspect encrypted traffic by SSL or TLS version:

Step 1 In the SSL policy where you want to control encrypted traffic by encryption protocol version, create a new SSL rule or edit an existing rule.

For detailed instructions, see Understanding and Creating SSL Rules, page 16-4.

Step 2 In the SSL rule editor, select the Version tab.

The Version tab appears.

Step 3 Select the protocol versions you want to match against: **SSL v3.0**, **TLS v1.0**, **TLS v1.1**, or **TLS v1.2**.

Step 4 Add or continue editing the rule.

You must apply the access control policy associated with the SSL policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.



Understanding Network Analysis and Intrusion Policies

Network analysis and intrusion policies work together as part of the ASA FirePOWER module intrusion detection and prevention feature. The term *intrusion detection* generally refers to the process of passively analyzing network traffic for potential intrusions and storing attack data for security analysis. The term *intrusion prevention* includes the concept of intrusion detection, but adds the ability to block or alter malicious traffic as it travels across your network.

In an intrusion prevention deployment, when the system examines packets:

- A **network analysis policy** governs how traffic is *decoded* and *preprocessed* so that it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt.
- An **intrusion policy** uses *intrusion and preprocessor rules* (sometimes referred to collectively as *intrusion rules*) to examine the decoded packets for attacks based on patterns. Intrusion policies are paired with *variable sets*, which allow you to use named values to accurately reflect your network environment.

Both network analysis and intrusion policies are invoked by a parent access control policy, but at different times. As the system analyzes traffic, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (additional preprocessing and intrusion rules) phase. Together, network analysis and intrusion policies provide broad and deep packet inspection. They can help you detect, alert on, and protect against network traffic that could threaten the availability, integrity, and confidentiality of hosts and their data.

The ASA FirePOWER moduleis delivered with several similarly named network analysis and intrusion policies (for example, Balanced Security and Connectivity) that complement and work with each other. By using system-provided policies, you can take advantage of the experience of the Cisco Vulnerability Research Team (VRT). For these policies, the VRT sets intrusion and preprocessor rule states, as well as provides the initial configurations for preprocessors and other advanced settings.

You can also create custom network analysis and intrusion policies. You can tune settings in custom policies to inspect traffic in the way that matters most to you.

You create, edit, save, and manage network analysis and intrusion policies using similar policy editors. When you are editing either type of policy, a navigation panel appears on the left side of the user interface; the right side displays various configuration pages.

This chapter contains a brief overview of the types of configurations the network analysis and intrusion policies govern, explains how the policies work together to examine traffic and generate records of policy violations, and provides basic information on navigating the policy editors. This chapter also explains the benefits and limitations of using custom versus system-provided policies. For more information, see the following sections:

- Understanding How Policies Examine Traffic For Intrusions, page 18-2
- Comparing System-Provided with Custom Policies, page 18-7
- Using the Navigation Panel, page 18-13
- Resolving Conflicts and Committing Policy Changes, page 18-15

To customize your intrusion deployment, see the following for your next steps:

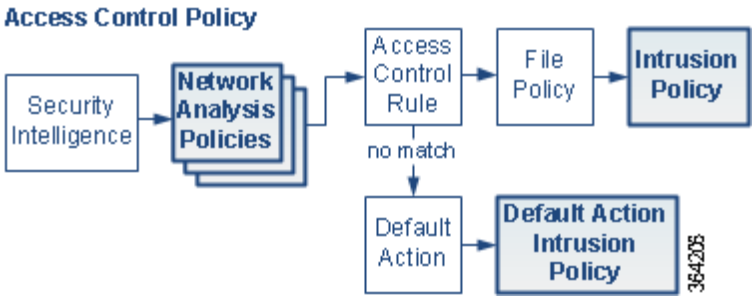
- Working with Variable Sets, page 2-13 explains how to configure the system's intrusion variables to accurately reflect your network environment. Even if you do not use custom policies, Cisco **strongly** recommends that you modify the default variables in the default variable set. Advanced users can create and use custom variable sets for pairing with one or more custom intrusion policies.
- Getting Started with Intrusion Policies, page 26-1 explains how to create and edit a simple custom intrusion policy.
- Controlling Traffic Using Intrusion and File Policies, page 11-1 explains how to configure the system to use intrusion policies to examine only the traffic you are interested in by associating intrusion policies with a parent access control policy. It also explains how to configure advanced intrusion policy performance options.
- Configuring Advanced Transport/Network Settings, page 24-1 explains how to configure advanced transport and network preprocessor settings that apply globally to all traffic. You configure these advanced settings in an access control policy rather than in a network analysis or intrusion policy.
- Getting Started with Network Analysis Policies, page 21-1 explains how to create and edit a simple custom network analysis policy.
- Customizing Preprocessing with Network Analysis Policies, page 20-2 explains how to change the default network analysis policy. For advanced users, this section also explains how to tailor preprocessing to specific security zones and networks by assigning custom network analysis policies to preprocess matching traffic.
- Using Layers in a Network Analysis or Intrusion Policy, page 19-1 explain how, in larger organizations or complex deployments, you can use building blocks called policy *layers* to more efficiently manage multiple network analysis or intrusion policies.

Understanding How Policies Examine Traffic For Intrusions

License: Protection

When the system analyzes traffic as part of your access control deployment, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (intrusion rules and advanced settings) phase.

The following diagram shows, in a simplified fashion, the order of traffic analysis in an inline, intrusion prevention and advanced malware protection (AMP) deployment. It illustrates how the access control policy invokes other policies to examine traffic, and in which order those policies are invoked. The network analysis and intrusion policy selection phases are highlighted.



In an inline deployment, the system can block traffic without further inspection at almost any step in the illustrated process. Security Intelligence, the SSL policy, network analysis policies, file policies, and intrusion policies can all either drop or modify traffic.

Similarly, at each step of the process, a packet could cause the system to generate an event. Intrusion and preprocessor events (sometimes referred to collectively as *intrusion events*) are indications that a packet or its contents may represent a security risk.

Note that for a single connection, although the system selects a network analysis policy before an access control rule as shown in the diagram, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.

For more information, see:

- Decoding, Normalizing, and Preprocessing: Network Analysis Policies, page 18-3
- Access Control Rules: Intrusion Policy Selection, page 18-4
- Intrusion Inspection: Intrusion Policies, Rules, and Variable Sets, page 18-5
- Intrusion Event Generation, page 18-6

Decoding, Normalizing, and Preprocessing: Network Analysis Policies

License: Protection

Without decoding and preprocessing, the system could not appropriately evaluate traffic for intrusions because protocol differences would make pattern matching impossible. As shown in the diagram in Understanding How Policies Examine Traffic For Intrusions, page 18-2, network analysis policies govern these traffic-handling tasks:

- **after** traffic is filtered by Security Intelligence
- **before** traffic can be inspected by file or intrusion policies

A network analysis policy governs packet processing in phases. First the system decodes packets through the first three TCP/IP layers, then continues with normalizing, preprocessing, and detecting protocol anomalies.

- The packet decoder converts packet headers and payloads into a format that can be easily used by the preprocessors and later, intrusion rules. Each layer of the TCP/IP stack is decoded in turn, beginning with the data link layer and continuing through the network and transport layers. The

packet decoder also detects various anomalous behaviors in packet headers. For more information, see Understanding Packet Decoding, page 24-16.

- In inline deployments, the inline normalization preprocessor reformats (normalizes) traffic to minimize the chances of attackers evading detection. It prepares packets for examination by other preprocessors and intrusion rules, and helps ensure that the packets the system processes are the same as the packets received by the hosts on your network. For more information, see Normalizing Inline Traffic, page 24-6.
- Various network and transport layers preprocessors detect attacks that exploit IP fragmentation, perform checksum validation, and perform TCP and UDP session preprocessing; see Configuring Transport & Network Layer Preprocessing, page 24-1.

  Note that some advanced transport and network preprocessor settings apply globally to all traffic handled by an access control policy. You configure these in the access control policy rather than in a network analysis policy; see Configuring Advanced Transport/Network Settings, page 24-1.
- Various application-layer protocol decoders normalize specific types of packet data into formats that the intrusion rules engine can analyze. Normalizing application-layer protocol encodings allows the system to effectively apply the same content-related intrusion rules to packets whose data is represented differently, and to obtain meaningful results. For more information, see Using Application Layer Preprocessors, page 22-1.
- The Modbus and DNP3 SCADA preprocessors detect traffic anomalies and provide data to intrusion rules. Supervisory Control and Data Acquisition (SCADA) protocols monitor, control, and acquire data from industrial, infrastructure, and facility processes such as manufacturing, production, water treatment, electric power distribution, airport and shipping systems, and so on. For more information, see Configuring SCADA Preprocessing, page 23-1.
- Several preprocessors allow you to detect specific threats, such as Back Orifice, portscans, SYN floods and other rate-based attacks; see Detecting Specific Threats, page 28-1.

  Note that you configure the sensitive data preprocessor, which detects sensitive data such as credit card numbers and Social Security numbers in ASCII text, in intrusion policies; see Detecting Sensitive Data, page 28-19.

In a newly created access control policy, one default network analysis policy governs preprocessing for *all* traffic for *all* intrusion policies invoked by the same parent access control policy. Initially, the system uses the Balanced Security and Connectivity network analysis policy as the default, but you can change it to another system-provided or custom network analysis policy. In a more complex deployment, advanced users can tailor traffic preprocessing options to specific security zones and networks by assigning different custom network analysis policies to preprocess matching traffic. For more information, see Comparing System-Provided with Custom Policies, page 18-7.

Access Control Rules: Intrusion Policy Selection

License: Protection

After initial preprocessing, access control rules (when present) evaluate traffic. In most cases, the first access control rule that a packet matches is the rule that handles that traffic; you can monitor, trust, block, or allow matching traffic.

When you allow traffic with an access control rule, the system can inspect the traffic for malware, prohibited files, and intrusions, in that order. Traffic not matching any access control rule is handled by the access control policy's default action, which can also inspect for intrusions.

Note All packets, **regardless** of which network analysis policy preprocesses them, are matched to configured access control rules—and thus are potentially subject to inspection by intrusion policies—in top-down order. For more information, see Limitations of Custom Policies, page 18-11.

The diagram in Understanding How Policies Examine Traffic For Intrusions, page 18-2 shows the flow of traffic through a device in an inline, intrusion prevention and AMP deployment, as follows:

- The access control rule allows matching traffic to proceed. The traffic is then inspected for prohibited files and malware by a file policy, and then for intrusions by an intrusion policy.
- In this scenario, the access control policy's default action allows matching traffic. The traffic is then inspected by an intrusion policy. You can (but do not have to) use a different intrusion policy when you associate intrusion policies with access control rules or the default action.

The example in the diagram does not include any blocking or trusting rules because the system does not inspect blocked or trusted traffic. For more information, see Using Rule Actions to Determine Traffic Handling and Inspection, page 6-7 and Setting Default Handling and Inspection for Network Traffic, page 4-4.

Intrusion Inspection: Intrusion Policies, Rules, and Variable Sets

License: Protection

You can use intrusion prevention as the system's last line of defense before traffic is allowed to proceed to its destination. Intrusion policies govern how the system inspects traffic for security violations and, in inline deployments, can block or alter malicious traffic. The main function of intrusion policies is to manage which intrusion and preprocessor rules are enabled and how they are configured.

Intrusion and Preprocessor Rules

An intrusion rule is a specified set of keywords and arguments that detects attempts to exploit vulnerabilities on your network; the system uses an intrusion rule to analyze network traffic to check if it matches the criteria in the rule. The system compares packets against the conditions specified in each rule and, if the packet data matches all the conditions specified in a rule, the rule triggers.

The system includes the following types of rules created by the VRT:

- *shared object intrusion rules*, which are compiled and cannot be modified (except for rule header information such as source and destination ports and IP addresses)
- *standard text intrusion rules*, which can be saved and modified as new custom instances of the rule.
- *preprocessor rules*, which are rules associated with preprocessors and packet decoder detection options in the network analysis policy. You cannot copy or edit preprocessor rules. Most preprocessor rules are disabled by default; you must enable them to use preprocessors to generate events and, in an inline deployment, drop offending packets.

When the system processes packets according to an intrusion policy, first a rule optimizer classifies all activated rules in subsets based on criteria such as: transport layer, application protocol, direction to or from the protected network, and so on. Then, the intrusion rules engine selects the appropriate rule subsets to apply to each packet. Finally, a multi-rule search engine performs three different types of searches to determine if the traffic matches the rule:

- The protocol field search looks for matches in particular fields in an application protocol.
- The generic content search looks for ASCII or binary byte matches in the packet payload.
- The packet anomaly search looks for packet headers and payloads that, rather than containing specific content, violate well-established protocols.

In a custom intrusion policy, you can tune detection by enabling and disabling rules, as well as by writing and adding your own standard text rules.

Variable Sets

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Most variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.

The system provides a single default variable set, which is comprised of predefined default variables. Most system-provided shared object rules and standard text rules use these predefined default variables to define networks and port numbers. For example, the majority of the rules use the variable `$HOME_NET` to specify the protected network and the variable `$EXTERNAL_NET` to specify the unprotected (or outside) network. In addition, specialized rules often use other predefined variables. For example, rules that detect exploits against web servers use the `$HTTP_SERVERS` and `$HTTP_PORTS` variables.



Tip Even if you use system-provided intrusion policies, Cisco **strongly** recommends you modify key default variables in the default set. When you use variables that accurately reflect your network environment, processing is optimized and the system can monitor relevant systems for suspicious activity. Advanced users can create and use custom variable sets for pairing with one or more custom intrusion policies. For more information, see Optimizing Predefined Default Variables, page 2-13.

Intrusion Event Generation

License: Protection

When the system identifies a possible intrusion, it generates an *intrusion* or *preprocessor event* (sometimes collectively called *intrusion events*). You can view the data to gain a greater understanding of the attacks against your network assets. In an inline deployment, the system can also drop or replace packets that you know to be harmful.

Each intrusion event includes an event header and contains information about the event name and classification; the source and destination IP addresses; ports; the process that generated the event; and the date and time of the event, as well as contextual information about the source of the attack and its target. For packet-based events, the system also logs a copy of the decoded packet header and payload for the packet or packets that triggered the event.

The packet decoder, the preprocessors, and the intrusion rules engine can all cause the system to generate an event. For example:

- If the packet decoder (configured in the network analysis policy) receives an IP packet that is less than 20 bytes, which is the size of an IP datagram without any options or payload, the decoder interprets this as anomalous traffic. If, later, the accompanying decoder rule in the intrusion policy that examines the packet is enabled, the system generates a preprocessor event.
- If the IP defragmentation preprocessor encounters a series of overlapping IP fragments, the preprocessor interprets this as a possible attack and, when the accompanying preprocessor rule is enabled, the system generates a preprocessor event.
- Within the intrusion rules engine, most standard text rules and shared object rules are written so that they generate intrusion events when triggered by packets.

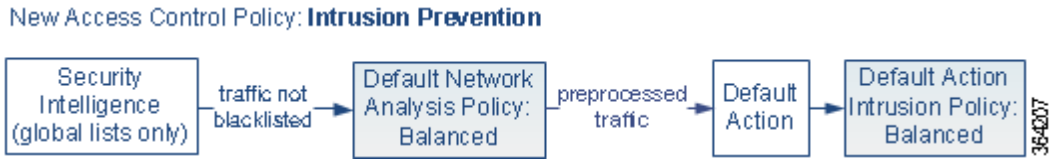
As the device accumulates intrusion events, you can begin your analysis of potential attacks. The system provides you with the tools you need to review intrusion events and evaluate whether they are important in the context of your network environment and your security policies.

Comparing System-Provided with Custom Policies

License: Protection

Creating a new access control policy is one of the first steps in managing traffic flow using the ASA FirePOWER module. By default, a newly created access control policy invokes system-provided network analysis and intrusion policies to examine traffic.

The following diagram shows how a newly created access control policy in an inline, intrusion-prevention deployment initially handles traffic. The preprocessing and intrusion prevention phases are highlighted.



Note how:

- A default network analysis policy governs the preprocessing of *all* traffic handled by the access control policy. Initially, the system-provided *Balanced Security and Connectivity network analysis policy* is the default.
- The default action of the access control policy allows all non-malicious traffic, as determined by the system-provided *Balanced Security and Connectivity intrusion policy.*
- The policy uses default Security Intelligence options (global whitelist and blacklist only), does not decrypt encrypted traffic with an SSL policy, and does not perform special handling and inspection of network traffic using access control rules.

A simple step you can take to tune your intrusion prevention deployment is to use a different set of system-provided network analysis and intrusion policies as your defaults. Cisco delivers several pairs of these policies with the ASA FirePOWER module.

Or, you can tailor your intrusion prevention deployment by creating and using custom policies. You may find that the preprocessor options, intrusion rule, and other advanced settings configured in those policies do not address the security needs of your network. By tuning your network analysis and intrusion policies you can configure, at a very granular level, how the system processes and inspects the traffic on your network for intrusions.

For more information, see:

- Understanding the System-Provided Policies, page 18-8
- Benefits of Custom Policies, page 18-9
- Limitations of Custom Policies, page 18-11

Understanding the System-Provided Policies

License: Protection

Cisco delivers several pairs of network analysis and intrusion policies with the ASA FirePOWER module. By using system-provided network analysis and intrusion policies, you can take advantage of the experience of the Cisco Vulnerability Research Team (VRT). For these policies, the VRT sets intrusion and preprocessor rule states, as well as provides the initial configurations for preprocessors and other advanced settings. You can use system-provided policies as-is, or you can use them as the base for custom policies.

Tip Even if you use system-provided network analysis and intrusion policies, you should configure the system's intrusion variables to accurately reflect your network environment. At a minimum, modify key default variables in the default set; see Optimizing Predefined Default Variables, page 2-13.

As new vulnerabilities become known, the VRT releases intrusion rule updates. These rule updates can modify any system-provided network analysis or intrusion policy, and can provide new and updated intrusion rules and preprocessor rules, modified states for existing rules, and modified default policy settings. Rule updates may also delete rules from system-provided policies and provide new rule categories, as well as modify the default variable set.

If a rule update affects your deployment, the system marks affected intrusion and network analysis policies as out of date, as well as their parent access control policies. You must reapply an updated policy for its changes to take effect.

For your convenience, you can configure rule updates to automatically reapply affected intrusion policies, either alone or in combination with affected access control policies. This allows you to easily and automatically keep your deployment up-to-date to protect against recently discovered exploits and intrusions.

To ensure up-to-date preprocessing settings, you **must** reapply access control policies, which also reapplies any associated SSL, network analysis, and file policies that are different from those currently running, and can also can update default values for advanced preprocessing and performance options. For more information, see Importing Rule Updates and Local Rule Files, page 46-9.

Cisco delivers the following network analysis and intrusion policies with the ASA FirePOWER module:

Balanced Security and Connectivity network analysis and intrusion policies

These policies are built for both speed and detection. Used together, they serve as a good starting point for most organizations. The system uses the Balanced Security and Connectivity policies and settings as defaults in most cases.

Connectivity Over Security network analysis and intrusion policies

These policies are built for organizations where connectivity (being able to get to all resources) takes precedence over network infrastructure security. The intrusion policy enables far fewer rules than those enabled in the Security over Connectivity policy. Only the most critical rules that block traffic are enabled.

Security Over Connectivity network analysis and intrusion policies

These policies are built for organizations where network infrastructure security takes precedence over user convenience. The intrusion policy enables numerous network anomaly intrusion rules that could alert on or drop legitimate traffic.

No Rules Active intrusion policy

In the No Rules Active intrusion policy, all intrusion rules and advanced settings are disabled. This policy provides a starting point if you want to create your own intrusion policy instead of basing it on the enabled rules in one of the other system-provided policies.

Caution Cisco uses another policy, Experimental Policy 1, for testing purposes. Do not use it unless instructed to do so by a Cisco representative.

Benefits of Custom Policies

License: Protection

You may find that the preprocessor options, intrusion rules, and other advanced settings configured in the system-provided network analysis and intrusion policies do not fully address the security needs of your organization.

Building custom policies can improve the performance of the system in your environment and can provide a focused view of the malicious traffic and policy violations occurring on your network. By creating and tuning custom policies you can configure, at a very granular level, how the system processes and inspects the traffic on your network for intrusions.

All custom policies have a base policy, also called a base layer, which defines the default settings for all configurations in the policy. A layer is a building block that you can use to efficiently manage multiple network analysis or intrusion policies; see Using Layers in a Network Analysis or Intrusion Policy, page 19-1.

In most cases, you base custom policies on system-provided policies, but you can use another custom policy. However, all custom policies have a system-provided policy as the eventual base in a policy chain. Because rule updates can modify system-provided policies, importing a rule update may affect you even if you are using a custom policy as your base. If a rule update affects your policies, the module interface marks affected policies as out of date. For more information, see Allowing Rule Updates to Modify a System-Provided Base Policy, page 19-4.

For more information, see:

- Benefits of a Custom Network Analysis Policy, page 18-9
- Benefits of Custom Intrusion Policies, page 18-10

Benefits of a Custom Network Analysis Policy

License: Protection

By default, one network analysis policy preprocesses all unencrypted traffic handled by the access control policy. That means that all packets are decoded and preprocessed according to the same settings, regardless of the intrusion policy (and therefore intrusion rule set) that later examines them.

Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default. A simple way to tune preprocessing is to create and use a custom network analysis policy as the default; see Setting the Default Network Analysis Policy for Access Control, page 20-3.

Tuning options available vary by preprocessor, but some of the ways you can tune preprocessors and decoders include:

- You can disable preprocessors that do not apply to the traffic you are monitoring. For example, the HTTP Inspect preprocessor normalizes HTTP traffic. If you are confident that your network does not include any web servers using Microsoft Internet Information Services (IIS), you can disable the preprocessor option that looks for IIS-specific traffic and thereby reduce system processing overhead.

  Note If you disable a preprocessor in a custom network analysis policy, but the system needs to use that preprocessor to later evaluate packets against an enabled intrusion or preprocessor rule, the system automatically enables and uses the preprocessor although the preprocessor remains disabled in the network analysis policy user interface.

- Specify ports, where appropriate, to focus the activity of certain preprocessors. For example, you can identify additional ports to monitor for DNS server responses or encrypted SSL sessions, or ports on which you decode telnet, HTTP, and RPC traffic

For advanced users with complex deployments, you can create multiple network analysis policies, each tailored to preprocess traffic differently. Then, you can configure the system to use those policies to govern the preprocessing of traffic using different security zones or networks.

Note Tailoring preprocessing using custom network analysis policies—especially multiple network analysis policies—is an advanced task. Because preprocessing and intrusion inspection are so closely related, you **must** be careful to allow the network analysis and intrusion policies examining a single packet to complement each other. For more information, see Limitations of Custom Policies, page 18-11.

Benefits of Custom Intrusion Policies

License: Protection

In a newly created access control policy initially configured to perform intrusion prevention, the default action allows all traffic, but first inspects it with the system-provided Balanced Security and Connectivity intrusion policy. Unless you add access control rules or change the default action, all traffic is inspected by that intrusion policy; see the diagram in Comparing System-Provided with Custom Policies, page 18-7.

To customize your intrusion prevention deployment, you can create multiple intrusion policies, each tailored to inspect traffic differently. Then, configure an access control policy with rules that specify which policy inspects which traffic. Access control rules can be simple or complex, matching and inspecting traffic using multiple criteria including security zone, network or geographical location, port, application, requested URL, or user. The scenario in Understanding How Policies Examine Traffic For Intrusions, page 18-2 shows a deployment where traffic is inspected by one of two intrusion policies.

The main function of intrusion policies is to manage which intrusion and preprocessor rules are enabled and how they are configured, as follows:

- Within each intrusion policy, you should verify that all rules applicable to your environment are enabled, and improve performance by disabling rules that are not applicable to your environment. In an inline deployment, you can specify which rules should drop or modify malicious packets. For more information, see Setting Rule States, page 27-19.
- You can modify existing rules and write new standard text rules as needed to catch new exploits or to enforce your security policies; see Understanding and Writing Intrusion Rules, page 30-1.

Other customizations you might make to an intrusion policy include:

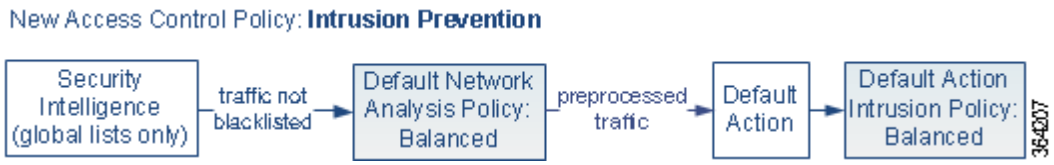
- The sensitive data preprocessor detects sensitive data such as credit card numbers and Social Security numbers in ASCII text. Note that other preprocessors that detect specific threats (back orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic) are configured in network analysis policies. For more information, see Detecting Specific Threats, page 28-1.
- Global thresholds cause the system to generate events based on how many times traffic matching an intrusion rule originates from or is targeted to a specific address or address range within a specified time period. This helps prevent the system from being overwhelmed with a large number of events. For more information, see Globally Limiting Intrusion Event Logging, page 29-1.
- Suppressing intrusion event notifications and setting thresholds for individual rules or entire intrusion policies can also can prevent the system from being overwhelmed with a large number of events. For more information, see Filtering Intrusion Event Notification Per Policy, page 27-21.
- In addition to intrusion events, you can enable logging to syslog facilities or send event data to an SNMP trap server. Per policy, you can specify intrusion event notification limits, set up intrusion event notification to external logging facilities, and configure external responses to intrusion events. For more information, see Configuring External Alerting for Intrusion Rules, page 39-1.

Limitations of Custom Policies

License: Protection

Because preprocessing and intrusion inspection are so closely related, you **must** be careful that your configuration allows the network analysis and intrusion policies processing and examining a single packet to complement each other.

By default, the system uses one network analysis policy to preprocess all traffic. The following diagram shows how a newly created access control policy in an inline, intrusion-prevention deployment initially handles traffic. The preprocessing and intrusion prevention phases are highlighted.



Notice how a default network analysis policy governs the preprocessing of all traffic handled by the access control policy. Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default.

A simple way to tune preprocessing is to create and use a custom network analysis policy as the default, as summarized in Benefits of a Custom Network Analysis Policy, page 18-9. However, if you disable a preprocessor in a custom network analysis policy but the system needs to evaluate preprocessed packets against an enabled intrusion or preprocessor rule, the system automatically enables and uses the preprocessor although it remains disabled in the network analysis policy user interface.

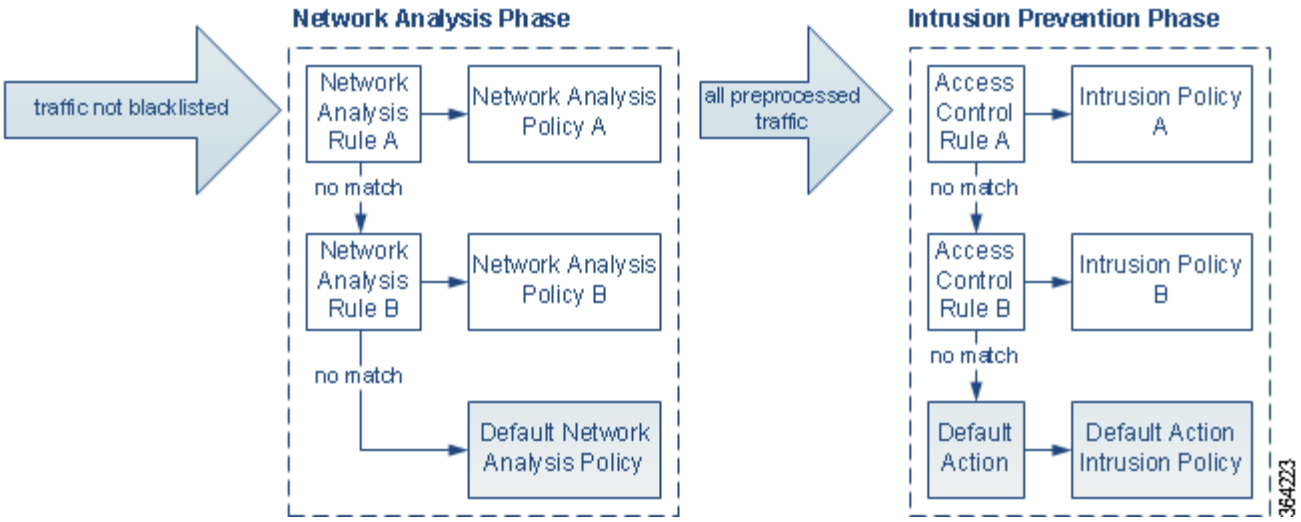
Note In order to get the performance benefits of disabling a preprocessor, you **must** make sure that none of your intrusion policies have enabled rules that require that preprocessor.

An additional challenge arises if you use multiple custom network analysis policies. For advanced users with complex deployments, you can tailor preprocessing to specific security zones and networks by assigning custom network analysis policies to preprocess matching traffic. To accomplish this, you add custom *network analysis rules* to your access control policy. Each rule has an associated network analysis policy that governs the preprocessing of traffic that matches the rule.

Tip You configure network analysis rules as an advanced setting in an access control policy. Unlike other types of rules in the ASA FirePOWER module, network analysis rules invoke—rather than being contained by—network analysis policies.

The system matches packets to any configured network analysis rules in top-down order by rule number. Traffic that does not match any network analysis rule is preprocessed by the default network analysis policy. While this allows you a great deal of flexibility in preprocessing traffic, keep in mind that all packets, **regardless** of which network analysis policy preprocessed them, are subsequently matched to access control rules—and thus to potential inspection by intrusion policies—in their own process. In other words, preprocessing a packet with a particular network analysis policy does **not** guarantee that the packet will be examined with any particular intrusion policy. You **must** carefully configure your access control policy so it invokes the correct network analysis and intrusion policies to evaluate a particular packet.

The following diagram shows in focused detail how the network analysis policy (preprocessing) selection phase occurs before and separately from the intrusion prevention (rules) phase. For simplicity, the diagram eliminates the file/malware inspection phases. It also highlights the default network analysis and default-action intrusion policies.



In this scenario, an access control policy is configured with two network analysis rules and a default network analysis policy:

- Network Analysis Rule A preprocesses matching traffic with Network Analysis Policy A. Later, you want this traffic to be inspected by Intrusion Policy A.
- Network Analysis Rule B preprocesses matching traffic with Network Analysis Policy B. Later, you want this traffic to be inspected by Intrusion Policy B.
- All remaining traffic is preprocessed with the default network analysis policy. Later, you want this traffic to be inspected by the intrusion policy associated with the access control policy's default action.

After the system preprocesses traffic, it can examine the traffic for intrusions. The diagram shows an access control policy with two access control rules and a default action:

- Access Control Rule A allows matching traffic. The traffic is then inspected by Intrusion Policy A.
- Access Control Rule B allows matching traffic. The traffic is then inspected by Intrusion Policy B.
- The access control policy's default action allows matching traffic. The traffic is then inspected by the default action's intrusion policy.

Each packet's handling is governed by a network analysis policy and intrusion policy pair, but the system does **not** coordinate the pair for you. Consider a scenario where you misconfigure your access control policy so that Network Analysis Rule A and Access Control Rule A do not process the same traffic. For example, you could intend the paired policies to govern the handling of traffic on a particular security zone, but you mistakenly use different zones in the two rules' conditions. This could cause traffic to be incorrectly preprocessed. For this reason, tailoring preprocessing using network analysis rules and custom policies is an **advanced** task.

Note that for a single connection, although the system selects a network analysis policy before an access control rule, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.

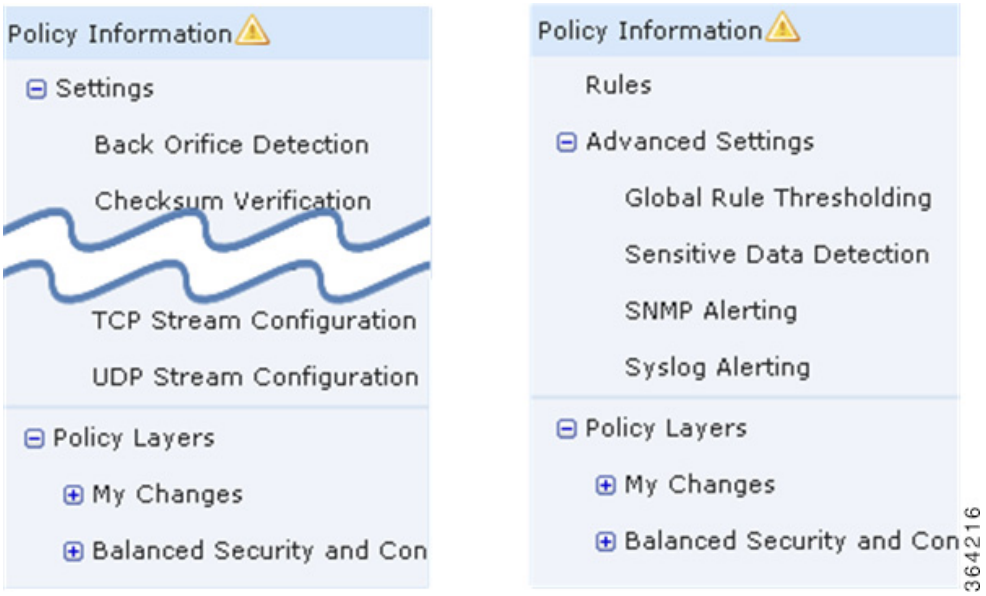
Using the Navigation Panel

License: Protection

Network analysis and intrusion policies use similar user interfaces to edit and save changes to their configurations; see:

- Editing Network Analysis Policies, page 21-3
- Editing Intrusion Policies, page 26-4

A navigation panel appears on the left side of the user interface when you are editing either type of policy. The following graphic shows the navigation panel for the network analysis policy (left) and the intrusion policy (right).



A dividing line separates the navigation panel into links to policy settings you can configure with (below) or without (above) direct interaction with policy layers. To navigate to any settings page, click its name in the navigation panel. Dark shading of an item in the navigation panel highlights your current settings page. For example, in the illustration above the Policy Information page would be displayed to the right of the navigation panel.

Policy Information

The Policy Information page provides configuration options for commonly used settings. As shown in the illustration for the network analysis policy panel above, a policy change icon (⚠) appears next to **Policy Information** in the navigation panel when the policy contains unsaved changes. The icon disappears when you save your changes.

Rules (intrusion policy only)

The Rules page in an intrusion policy allows you to configure rule states and other settings for shared object rules, standard text rules, and preprocessor rules. For more information, see Tuning Intrusion Policies Using Rules, page 27-1.

Settings (network analysis policy) and Advanced Settings (intrusion policy)

The Settings page in a network analysis policy allows you to enable or disable preprocessors and access preprocessor configuration pages. Expanding the **Settings** link displays sublinks to individual configuration pages for all enabled preprocessors in the policy. For more information, see Configuring Preprocessors in a Network Analysis Policy, page 21-6.

The Advanced Settings page in an intrusion policy allows you to enable or disable advanced settings and access configuration pages for those advanced settings. Expanding the **Advanced Settings** link displays sublinks to individual configuration pages for all enabled advanced settings in the policy. For more information, see Configuring Advanced Settings in an Intrusion Policy, page 26-6.

Policy Layers

The Policy Layers page displays a summary of the layers that comprise your network analysis or intrusion policy. Expanding the Policy Layers link displays sublinks to summary pages for the layers in your policy. Expanding each layer sublink displays further sublinks to the configuration pages for all rules, preprocessors, or advanced settings that are enabled in the layer. For more information, see Using Layers in a Network Analysis or Intrusion Policy, page 19-1.

Resolving Conflicts and Committing Policy Changes

License: Protection

When you edit a network analysis or intrusion policy, you must save (or *commit*) your changes before the system recognizes them.

Note After you save, you must apply a network analysis or intrusion policy for your changes to take effect. If you apply a policy without saving, the system uses the most recently saved configuration. Although you can reapply an intrusion policy independently, network analysis policies are applied with their parent access control policy.

Resolving Editing Conflicts

The Network Analysis Policy page and Intrusion Policy page display whether each policy has unsaved changes; see Editing Network Analysis Policies, page 21-3 and Editing Intrusion Policies, page 26-4.

Cisco recommends that only one person edit a policy at a time. If you are editing the same network analysis or intrusion policy via multiple user interface instances as the same user, and you save your changes for one instance, you cannot save your changes for the other instance.

Resolving Configuration Dependencies

To perform their particular analysis, many preprocessors and intrusion rules require that traffic first be decoded or preprocessed in a certain way, or have other dependencies. When you save a network analysis or intrusion policy, the system either automatically enables required settings, or warns you that disabled settings will have no effect on traffic, as follows:

- You cannot save an intrusion policy if you added an SNMP rule alert but did not configure SNMP alerting. You must either configure SNMP alerting or disable the rule alert, then save again.
- You cannot save an intrusion policy if it includes enabled sensitive data rules but you have not enabled the sensitive data preprocessor. You must either allow the system to enable the preprocessor and save the policy, or disable the rules and save again.
- If you disable a required preprocessor in a network analysis policy, you can still save the policy. However, the system automatically uses the disabled preprocessor with its current settings, even though the preprocessor remains disabled in the user interface. For more information, see Limitations of Custom Policies, page 18-11.
- If you disable inline mode in a network analysis policy but enable the Inline Normalization preprocessor, you can still save the policy. However, the system warns you that normalization settings will be ignored. Disabling inline mode also causes the system to ignore other settings that allow preprocessors to modify or block traffic, including checksum verification and rate-based attack prevention. For more information, see Allowing Preprocessors to Affect Traffic in Inline Deployments, page 21-5 and Normalizing Inline Traffic, page 24-6.

Committing, Discarding, and Caching Policy Changes

While editing a network analysis or intrusion policy, if you exit the policy editor without saving your changes, the system caches those changes. Your changes are cached even when you log out of the system or experience a system crash. The system cache can store unsaved changes for one network analysis and one intrusion policy; you must commit or discard your changes before editing another policy of the same type. The system discards the cached changes when you edit another policy without saving your changes to the first policy, or when you import an intrusion rule update.

You can commit or discard policy changes on the Policy Information page of either the network analysis or intrusion policy editor; see Editing Network Analysis Policies, page 21-3 and Editing Intrusion Policies, page 26-4.

The following table summarizes how to save or discard changes to a network analysis or intrusion policy.

Table 18-1 Committing Changes to a Network Analysis or Intrusion Policy

To...	On the Policy Information page, you can...
save changes to the policy	click **Commit Changes**. Optionally, enter a comment; click **OK** to continue committing.
discard all unsaved changes	click **Discard Changes**, then click **OK** to discard your changes and go to the Intrusion Policy page. If you do not want to discard your changes, click **Cancel** to return to the Policy Information page.
exit the policy, but cache changes	select any menu or other path to another page. On exiting, click **Leave page** when prompted, or click **Stay on page** to remain in the advanced editor.



Using Layers in a Network Analysis or Intrusion Policy

Larger organizations with many ASA FirePOWER modules may have many intrusion policies and network analysis policies to support the unique needs of different departments, business units or, in some instances, different companies. Configurations in both policy types are contained in building blocks called *layers,* which you can use to efficiently manage multiple policies.

Layers in intrusion and network analysis policies work in essentially the same way. You can create and edit either policy type without consciously using layers. You can modify your policy configurations and, if you have not added user layers to your policy, the system automatically includes your changes in a single configurable layer that is initially named *My Changes*. Optionally, you can also add up to 200 layers where you can configure any combination of settings. You can copy, merge, move, and delete user layers and, most important, share individual user layers with other policies of the same type.

See the following sections for more information:

- Understanding the Layer Stack, page 19-1 describes the user-configurable and built-in layers that comprise a basic policy.
- Managing Layers, page 19-5 explains how to use layers in your policies.

Understanding the Layer Stack

License: Protection

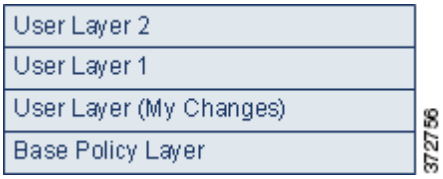
A network analysis or intrusion policy where you do not add layers includes the built-in, read-only base policy layer and a single user-configurable layer that is initially named My Changes. You can copy, merge, move, or delete any user-configurable layer and set any user-configurable layer to be shared by other policies of the same type.

Each policy layer contains complete configurations for either all preprocessors in a network analysis policy or all intrusion rules and advanced settings in an intrusion policy. The lowest, base policy layer includes all the settings from the base policy you selected when you created the policy. A setting in a higher layer takes precedence over the same setting in a lower layer. Features not explicitly set in a layer *inherit* their settings from the next highest layer where they are explicitly set.

The system *flattens* the layers, that is, it applies only the cumulative effect of all settings, when it handles network traffic.

Tip You can create an intrusion or network analysis policy based solely on the default settings in the base policy.

The following figure shows an example layer stack that, in addition to the base policy layer and the initial My Changes layer, also includes two additional user-configurable layers, *User Layer 1* and *User Layer 2*. Note in the figure that each user-configurable layer that you add is initially positioned as the highest layer in the stack; thus, User Layer 2 in the figure was added last and is highest in the stack.



Note the following points when working with multiple layers:

- When the highest layer in your policy is a read-only layer, or a shared layer as described in Sharing Layers Between Policies, page 19-9, the system automatically adds a user-configurable layer as the highest layer in your intrusion policy if you do either of the following:
  - modify a rule action (that is, a rule state, event filtering, dynamic state, or alerting) on the intrusion policy Rules page. See Tuning Intrusion Policies Using Rules, page 27-1 for more information.
  - enable, disable, or modify any preprocessor, intrusion rule, or advanced setting.

  All settings in the system-added layer are inherited except for the changes that resulted in the new layer.
- When the highest layer is a shared layer, the system adds a layer when you take either of the following actions:
  - share the highest layer with other policies
  - add a shared layer to your policy
- Regardless of whether you allow rule updates to modify your policy, changes in a rule update never override changes you make in a layer. This is because changes in a rule update are made in the base policy, which determines the default settings in your base policy layer; your changes are always made in a higher layer, so they override any changes that a rule update makes to your base policy. See Importing Rule Updates and Local Rule Files, page 46-9 for more information.

See Understanding the Base Layer, page 19-2 for more information.

Understanding the Base Layer

License: Protection

The base layer, also referred to as the base policy, of an intrusion or network analysis policy defines the default settings for all configurations in the policy, and is the lowest layer in the policy. When you create a new policy and change a setting without adding new layers, the change is stored in the My Changes layer, and overrides—but does not change—the setting in the base policy.

See the following sections for more information:

- Understanding System-Provided Base Policies, page 19-3
- Understanding Custom Base Policies, page 19-3

- Changing the Base Policy, page 19-3
- Allowing Rule Updates to Modify a System-Provided Base Policy, page 19-4

Understanding System-Provided Base Policies

License: Protection

Cisco delivers several pairs of network analysis and intrusion policies with the ASA FirePOWER module. By using system-provided network analysis and intrusion policies, you can take advantage of the experience of the Cisco Vulnerability Research Team (VRT). For these policies, the VRT sets intrusion and preprocessor rule states, as well as provides the initial configurations for preprocessors and other advanced settings. You can use these system-provided policies as-is, or you can use them as the base for custom policies.

If you use a system-provided policy as your base, importing rule updates may modify settings in your base policy. However, you can configure a custom policy to not automatically make these changes to its system-provided base policy. This allows you to update system-provided base policies manually, on a schedule independent of rule update imports. In either case, changes that a rule update makes to your base policy do not change or override settings in your My Changes or any other layer. For more information, see Allowing Rule Updates to Modify a System-Provided Base Policy, page 19-4.

System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. For more information, see Understanding the System-Provided Policies, page 18-8.

Understanding Custom Base Policies

License: Protection

If you do not want to use a system-provided policy as the base policy in your network analysis or intrusion policy, you can use a custom policy as your base. You can tune settings in custom policies to inspect traffic in ways that matter most to you so you can improve both the performance of your device and your ability to respond effectively to the events they generate.

You can chain up to five custom policies, with four of the five using one of the other four previously created policies as its base policy; the fifth must use a system-provided policy as its base.

Changes you make to a custom policy that you use as the base for another policy are automatically used as the default settings the of policy that uses the base. Additionally, because all policies have a system-provided policy as the eventual base in a policy chain, importing a rule update may affect your policy even if you use a custom base policy. If the first custom policy in a chain (the one that uses the system-provided policy as its base) allows rule updates to modify its base policy, your policy may be affected. For information on changing this setting, see Allowing Rule Updates to Modify a System-Provided Base Policy, page 19-4.

Regardless of how they are made, changes to your base policy—whether by a rule update or when you modify a custom policy that you use as a base policy—do not change or override settings in your My Changes or any other layer.

Changing the Base Policy

License: Protection

You can select a different base policy for your network analysis or intrusion policy and, optionally, allow rule updates to modify a system-provided base policy, without affecting modifications in higher layers.

To change the base policy:

Step 1 While editing your policy, click **Policy Information** in the navigation panel.

The Policy Information page appears.

Step 2 Select a base policy from the **Base Policy** drop-down list.

Step 3 Optionally, if you choose a system-provided base policy, click **Manage Base Policy** to specify whether intrusion rule updates can automatically modify your base policy.

For more information, see Allowing Rule Updates to Modify a System-Provided Base Policy, page 19-4.

Step 4 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Allowing Rule Updates to Modify a System-Provided Base Policy

License: Protection

Rule updates that you import provide system-provided policies with modified network analysis preprocessor settings, modified intrusion policy advanced settings, new and updated intrusion rules, and modified states for existing rules. Rule updates can also delete rules and provide new rule categories and default variables. See Importing Rule Updates and Local Rule Files, page 46-9 for more information.

Rule updates always modify system-provided policies with any changes to preprocessors, advanced settings, and rules. Changes to default variables and rule categories are handled at the system level. See Understanding System-Provided Base Policies, page 19-3 for more information.

When you use a system-provided policy as your base policy, you can allow rule updates to modify your base policy which, in this case, is a copy of the system-provided policy. If you allow rule updates to update your base policy, a new rule update makes the same changes in your base policy that it makes to the system-provided policy that you use as your base policy. If you have not modified the corresponding setting, a setting in your base policy determines the setting in your policy. However, rule updates do not override changes you make in your policy.

If you do not allow rule updates to update your base policy, you can manually update your base policy after importing one or more rule updates.

Rule updates always delete intrusion rules that VRT deletes, regardless of the rule state in your intrusion policy or whether you allow rule updates to update your base intrusion policy. Until you reapply your changes to network traffic, rules in your currently applied intrusion policies behave as follows:

- Disabled rules remain disabled.
- Rules set to Generate Events continue to generate events when triggered.
- Rules set to Drop and Generate Events continue to generate events and drop offending packets when triggered.

Rule updates do not modify a custom base policy unless both of the following conditions are met:

- You allow rule updates to modify the system-provided base policy of the parent policy, that is, the policy that originated the custom base policy.
- You have not made changes in the parent policy that override the corresponding settings in the parent's base policy.

When both conditions are met, changes in the rule update are passed to the child policy, that is, the policy using the custom base policy, when you save the parent policy.

For example, if a rule update enables a previously disabled intrusion rule, and you have not modified the rule's state in the parent intrusion policy, the modified rule state is passed to the base policy when you save the parent policy.

Likewise, if a rule update modifies a default preprocessor setting and you have not modified the setting in the parent network analysis policy, the modified setting is passed to the base policy when you save the parent policy.

See Changing the Base Policy, page 19-3 for more information.

To allow rule updates to modify a system-provided base policy:

Step 1 While editing a policy that uses a system-provided policy as its base policy, click **Policy Information** in the navigation panel.

The Policy Information page appears.

Step 2 Click **Manage Base Policy**.

The Base Policy summary page appears.

Step 3 Select or clear the **Update when a new Rule Update is installed** check box.

When you save your policy with the check box cleared and then import a rule update, an **Update Now** button appears on the Base Policy summary page and the status message on the page updates to inform you that the policy is out of date. Optionally, you can click **Update Now** to update your base policy with the changes in the most recently imported rule update.

Step 4 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Managing Layers

License: Protection

The Policy Layers page provides a single-page summary of the complete layer stack for your network analysis or intrusion policy. On this page you can add shared and unshared layers, copy, merge, move, and delete layers, access the summary page for each layer, and access configuration pages for enabled, disabled, and overridden configurations within each layer.

For each layer, you can view the following information:

- whether the layer is a built-in, shared user, or unshared user layer
- which layers contain the highest, that is the effective, preprocessor or advanced setting configurations, by feature name
- in an intrusion policy, the number of intrusion rules whose states are set in the layer, and the number of rules set to each rule state.

The feature name in the summary for each layer indicates which configurations are enabled, disabled, overridden, or inherited in the layer, as follows:

When the feature is...	The feature name is...
enabled in the layer	written in plain text
disabled in the layer	struck out
overridden by the configuration in a higher layer	written in italic text
inherited from a lower layer	not present

This page also provides a summary of the net effect of all enabled preprocessors (network analysis) or advanced settings (intrusion) and, for intrusion policies, intrusion rules.

The following table lists the actions available on the Policy Layers page.

Table 19-1 Network Analysis and Intrusion Policy Layer Configuration Actions

To...	You can...
display the Policy Information page	click **Policy Summary**. See Tuning Intrusion Policies Using Rules, page 27-1, Getting Started with Network Analysis Policies, page 21-1, and Getting Started with Intrusion Policies, page 26-1 for information on the actions you can take on the Policy Information page.
display the summary page for a layer	click the layer name in the row for the layer or, alternately, click the edit icon ( ) next to a user layer. You can also click the view icon ( ) to access the read-only summary page for a shared layer. See Sharing Layers Between Policies, page 19-9, Configuring Preprocessors and Advanced Settings in Layers, page 19-14, and Configuring Intrusion Rules in Layers, page 19-11 for information on actions you can take on the summary page for a layer.
access a layer-level preprocessor or advanced setting configuration page	click the feature name in the row for the layer. Note that configuration pages are read-only in the base policy and in shared layers. See Configuring Preprocessors and Advanced Settings in Layers, page 19-14 for more information.
access a layer-level rule configuration page filtered by rule state type	click the icon for drop and generate events ( ), generate events ( ), or disabled ( ) in the summary for the layer. No rules are displayed if the layer contains no rules set to the selected rule state.
add a layer to your policy	see Adding a Layer, page 19-7.
add a shared layer from another policy	see Sharing Layers Between Policies, page 19-9.
change a layer's name or description	see Changing a Layer's Name and Description, page 19-7.
move, copy, or delete a layer	see Moving, Copying, and Deleting Layers, page 19-8.
merge a layer into the next layer beneath it	see Merging Layers, page 19-8.

To use the Policy Layers page:

Step 1 While editing your policy, click **Policy Layers** in the navigation panel.

The Policy Layers summary page appears.

Step 2 You can take any of the actions in the Network Analysis and Intrusion Policy Layer Configuration Actions table.

Step 3 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Adding a Layer

License: Protection

You can add up to 200 layers to a network analysis or intrusion policy. When you add a layer, it appears as the highest layer in your policy. The initial state is Inherit for all features and, in an intrusion policy, no event filtering, dynamic state, or alerting rule actions are set.

To add a layer to your network analysis or intrusion policy:

Step 1 While editing your policy, click **Policy Layers** in the navigation panel.

The Policy Layers page appears.

Step 2 Click the add layer icon ( ) next to User Layers.

The Add Layer pop-up window appears.

Step 3 Type a unique layer **Name** and click **OK**.

The new layer appears as the topmost layer under User Layers.

Step 4 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Changing a Layer's Name and Description

License: Protection

You can change the name of a user-configurable layer in your network analysis or intrusion policy and, optionally, add or modify a description that is visible when you edit the layer.

To change a layer's name and add or modify its description:

Step 1 While editing your policy, click **Policy Layers** in the navigation panel.

The Policy Layers page appears.

Step 2 Click the edit icon ( ) next to the user layer you want to edit.

The summary page for the layer appears.

Step 3 You can take the following actions:

- Modify the layer **Name**.

- Add or modify the layer **Description**.

Step 4 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Moving, Copying, and Deleting Layers

License: Protection

You can copy, move, or delete a user layer in your network analysis or intrusion policy, including the initial My Changes layer. Note the following considerations:

- When you copy a layer, the copy appears as the highest layer.
- Copying a shared layer creates an unshared copy which, optionally, you can then share with other policies.
- You cannot delete a shared layer; a layer with sharing enabled that you have not shared with another policy is not a shared layer.

To copy, move, or delete a layer:

Step 1 While editing your policy, click **Policy Layers** in the navigation panel.

The Policy Layers page appears.

Step 2 You can take the following actions:

- To copy a layer, click the copy icon () for the layer you want to copy.

  The page refreshes and a copy of the layer appears as the highest layer.
- To move a layer up or down within the User Layers page area, click any open area in the layer summary and drag until the position arrow () points to a line above or below a layer where you want to move the layer.

  The screen refreshes and the layer appears in the new location.
- To delete a layer, click the delete icon () for the layer you want to delete, then click **OK**

  The page refreshes and the layer is deleted.

Step 3 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Merging Layers

License: Protection

You can merge a user-configurable layer in your network analysis or intrusion policy with the next user layer beneath it. A merged layer retains all settings that were unique to either layer, and accepts the settings from the higher layer if both layers included settings for the same preprocessor, intrusion rule, or advanced setting. The merged layer retains the name of the lower layer.

In the policy where you create a shared layer that you add to other policies, you can merge an unshared layer immediately above the shared layer with the shared layer, but you cannot merge the shared layer with an unshared layer beneath it.

In a policy where you add a shared layer that you created in another policy, you can merge the shared layer into an unshared layer immediately beneath it and the resulting layer is no longer shared; you cannot merge an unshared layer into a shared layer beneath it.

To merge a user layer with a user layer beneath it:

Step 1 While editing your policy, click **Policy Layers** in the navigation panel.

The Policy Layers page appears.

Step 2 Click the merge icon ( ) in the upper of the two layers, then click **OK**.

The page refreshes and the layer is merged with the layer beneath it.

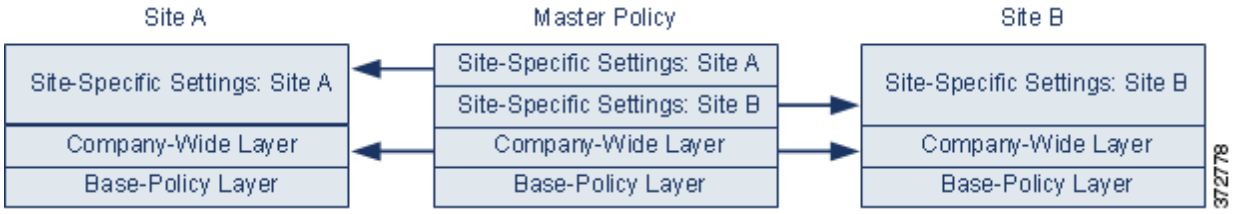
Step 3 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Sharing Layers Between Policies

License: Protection

You can share a user-configurable layer with other policies of the same type (intrusion or network analysis). When you modify a configuration within a shared layer and then commit your changes, the system updates all policies that use the shared layer and provides you with a list of all affected policies. You can only modify shared layer feature configurations in the policy where you created the layer.

The following figure shows an example master policy that serves as the source for site-specific policies.



The master policy in the figure includes a company-wide layer with settings applicable to the policies at Site A and Site B. It also includes site-specific layers for each policy. For example, in the case of a network analysis policy Site A might not have web servers on the monitored network and would not require the protection or processing overhead of the HTTP Inspect preprocessor, but both sites would likely require TCP stream preprocessing. You could enable TCP stream processing in the company-wide layer that you share with both sites, disable the HTTP Inspect preprocessor in the site-specific layer that you share with Site A, and enable the HTTP Inspect preprocessor in the site-specific layer that you share with Site B. By editing configurations in a higher layer in the site-specific policies, you could also further tune the policy for each site if necessary with any configuration adjustments.

It is unlikely that the flattened net settings in the example master policy would be useful for monitoring traffic, but the time saved in configuring and updating the site-specific policies makes this a useful application of policy layers.

Many other layer configurations are possible. For example, you could define policy layers by company, by department, or by network. In the case of an intrusion policy, you could also include advanced settings in one layer and rule settings in another.

Tip You cannot add a shared layer to a policy when your base policy is a custom policy where the layer you want to share was created. When you attempt to save your changes, an error message indicates that the policy includes a circular dependency. See Understanding Custom Base Policies, page 19-3 for more information.

To share a layer with other policies, you must do the following:

- Enable sharing on the layer summary page of the layer you want to share.
- Add the shared layer on the Policy Layers page of the policy where you want to share it.

You cannot disable sharing for a layer that is in use in another policy; you must first delete the layer from the other policy or delete the other policy.

To enable or disable sharing a layer with other policies:

Step 1 While editing your policy, click **Policy Layers** in the navigation panel.

The Policy Layers page appears.

Step 2 Click the edit icon ( ) next to the layer you want to share with other policies.

The summary page for the layer appears.

Step 3 Select (enable) or clear (disable) the **Sharing** check box.

Step 4 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

To add a shared layer to your policy:

Step 1 While editing your policy, click **Policy Layers** in the navigation panel.

The Policy Layers page appears.

Step 2 Click the add shared layer icon ( ) next to User Layers.

The Add Shared Layer pop-up window appears.

Step 3 Select the shared layer you want to add from the Add Shared Layer drop-down list, then click **OK**.

The Policy Layers summary page appears and the shared layer you selected appears as the highest layer in your policy.

If there are no shared layers in any other policies, no drop-down list appears; click **OK** or **Cancel** on the pop-up window to return to the Policy Layers summary page.

Step 4 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Configuring Intrusion Rules in Layers

License: Protection

In an intrusion policy, you can set the rule state, event filtering, dynamic state, alerting, and rule comments for a rule in any user-configurable layer. After accessing the layer where you want to make your changes, you add settings on the Rules page for the layer the same as you would on the intrusion policy Rules page; see Tuning Intrusion Policies Using Rules, page 27-1.

You can view individual layer settings on the Rules page for the layer, or view the net effect of all settings on the policy view of the Rules page. When you modify rule settings on the policy view of the Rules page, you are modifying the highest user-configurable layer in the policy. You can switch to another layer using the layer drop-down list on any Rules page.

The following table describes the effects of configuring the same type of setting in multiple layers.

Table 19-2 Layer Rule Settings

You can set...	Of this setting type...	To...
one	rule state	override a rule state set for the rule in a lower layer, and ignore all thresholds, suppressions, rate-based rule states, and alerts for that rule configured in lower layers. See Setting Rule States, page 27-19 for more information. If you want a rule to inherit its state from the base policy or a lower layer, set the rule state to Inherit. Note that when you are working on the intrusion policy Rules page, you cannot set a rule state to Inherit. Note also that rule state settings are color-coded when you view them on the Rules page for a specific layer: rules whose effective state is set in a lower layer are highlighted in yellow; rules whose effective state is set in a higher layer are highlighted in red; rules whose effective state is set in the current layer are not highlighted. Because the intrusion policy Rules page is a composite view of the net effect of all rule settings, rule states are not color-coded on this page.
one	threshold SNMP alert	override a setting of the same type for the rule in a lower layer. Note that setting a threshold overwrites any existing threshold for the rule in the layer. See Configuring Event Thresholding, page 27-21 and Adding SNMP Alerts, page 27-31 for more information.
one or more	suppression rate-based rule state	cumulatively combine settings of the same type for each selected rule down to the first layer where a rule state is set for the rule. Settings below the layer where a rule state is set are ignored. See Configuring Suppression Per Intrusion Policy, page 27-25 and Adding Dynamic Rule States, page 27-28 for more information.
one or more	comment	add a comment to a rule. Comments are rule-specific, not policy- or layer-specific. You can add one or more comments to a rule in any layer. See Adding a Rule Comment for a Rule, page 27-8 for more information.

For example, if you set a rule state to Drop and Generate Events in one layer and to Disabled in a higher layer, the intrusion policy Rules page shows that the rule is disabled.

In another example, if you set a source-based suppression for a rule to 192.168.1.1 in one layer, and you also set a destination-based suppression for the rule to 192.168.1.2 in another layer, the Rules page shows that the cumulative effect is to suppress events for the source address 192.168.1.1 and the

destination address 192.168.1.2. Note that suppression and rate-based rule state settings cumulatively combine settings of the same type for each selected rule down to the first layer where a rule state is set for the rule. Settings below the layer where a rule state is set are ignored.

To modify rules in a layer:

Step 1 While editing your intrusion policy, expand **Policy Layers** in the navigation panel and expand the policy layer you want to modify.

Step 2 Click **Rules** immediately beneath the policy layer you want to modify.

The Rules page for the layer appears.

You can modify any of the settings in the Layer Rule Settings table. See Tuning Intrusion Policies Using Rules, page 27-1 for more information on configuring intrusion rules.

To delete an individual setting from an editable layer, double-click the rule message on the Rules page for the layer to display rule details. Click **Delete** next to the setting you want to delete, then click **OK** twice.

Step 3 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Removing Multi-Layer Rule Settings

License: Protection

You can select one or more rules on the intrusion policy Rules page and then simultaneously remove a specific type of event filter, dynamic state, or alerting from multiple layers in your intrusion policy.

The system removes the setting type downward through each layer where it is set until it removes all the settings or encounters a layer where a rule state is set for the rule. If it encounters a layer where a rule state is set, it removes the setting from that layer and stops removing the setting type.

When the system encounters the setting type in a shared layer or in the base policy, and if the highest layer in the policy is editable, the system copies the remaining settings and rule state for the rule to that editable layer. Otherwise, if the highest layer in the policy is a shared layer, the system creates a new editable layer above the shared layer and copies the remaining settings and rule state for the rule to that editable layer.

Note Removing rule settings derived from a shared layer or the base policy causes any changes to this rule from lower layers or the base policy to be ignored. To stop ignoring changes from lower layers or the base policy, set the rule state to **Inherit** on the summary page for the topmost layer. See Setting Rule States, page 27-19 for more information.

To remove rule settings in multiple layers:

Step 1 While editing your intrusion policy, click **Rules** immediately beneath Policy Information in the navigation panel.

Tip You can also select **Policy** from the layer drop-down list on the Rules page for any layer, or select **Manage Rules** on the Policy Information page.

The intrusion policy Rules page appears.

Step 2 Select the rule or rules from which you want to remove multiple settings. You have the following options:

- To select a specific rule, select the check box next to the rule.
- To select all the rules in the current list, select the check box at the top of the column.

See Understanding Rule Filtering in an Intrusion Policy, page 27-9 and Setting a Rule Filter in an Intrusion Policy, page 27-17 for information on locating rules.

Step 3 You have the following options:

- To remove all thresholds for a rule, select **Event Filtering > Remove Thresholds**.
- To remove all suppression for a rule, select **Event Filtering > Remove Suppressions**.
- To remove all rate-based rule states for a rule, select **Dynamic State > Remove Rate-Based Rule States**.
- To remove all SNMP alert settings for a rule, select **Alerting > Remove SNMP Alerts**.

A confirmation pop-up window appears.

Note Removing rule settings derived from a shared layer or the base policy causes any changes to this rule from lower layers or the base policy to be ignored. To stop ignoring changes from lower layers or the base policy, set the rule state to **Inherit** on the summary page for the topmost layer. See Setting Rule States, page 27-19 for more information.

Step 4 Click **OK**.

The system removes the selected setting and copies the remaining settings for the rule to the highest editable layer in the policy. See the introduction to this procedure for conditions that affect how the system copies the remaining settings.

Step 5 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Accepting Rule Changes from a Custom Base Policy

License: Protection

When a custom network analysis or intrusion policy where you have not added layers uses another custom policy as its base policy, you must set a rule to inherit its rule state if:

- you delete an event filter, dynamic state, or SNMP alert set for the rule in the base policy, and
- you want the rule to accept subsequent changes that you make to it in the other custom policy that you use as your base policy

The following procedure explains how to accomplish this. See Removing Multi-Layer Rule Settings, page 19-12 to accept settings for these rules in a policy where you have added layers.

To accept rule changes in a policy where you have not added layers:

Step 1 While editing your intrusion policy, expand the **Policy Layers** link in the navigation panel, then expand the **My Changes** link.

Step 2 Click the **Rules** link immediately beneath My Changes.

The Rules page for the My Changes layer appears.

Step 3 Select the rule or rules whose settings you want to accept. You have the following options:

- To select a specific rule, select the check box next to the rule.
- To select all the rules in the current list, select the check box at the top of the column.

See Understanding Rule Filtering in an Intrusion Policy, page 27-9 and Setting a Rule Filter in an Intrusion Policy, page 27-17 for information on locating rules.

Step 4 Select **Inherit** from the **Rule State** drop-down list.

Step 5 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Configuring Preprocessors and Advanced Settings in Layers

License: Protection

You use similar mechanisms to configure preprocessors in a network analysis policy and advanced settings in an intrusion policy. You can enable and disable preprocessors on the network analysis Settings page and intrusion policy advanced settings on the intrusion policy Advanced Settings page. These pages also provide summaries of the effective states for all relevant features. For example, if the network analysis SSL preprocessor is disabled in one layer and enabled in a higher layer, the Settings page shows it as enabled. Changes made on these pages appear in the top layer of the policy.

You can also enable or disable preprocessors or advanced settings and access their configuration pages on the summary page for a user-configurable layer. On this page you can modify the layer name and description and configure whether to share the layer with other policies of the same type; see Sharing Layers Between Policies, page 19-9 for more information. You can switch to the summary page for another layer by selecting the layer name beneath **Policy Layers** in the navigation panel.

When you enable a preprocessor or advanced setting, a sublink to the configuration page for that feature appears beneath the layer name in the navigation panel, and an edit icon ( ) appears next to the feature on the summary page for the layer; these disappear when you disable the feature in the layer or set it to Inherit.

Setting the state (enabled or disabled) for a preprocessor or advanced setting overrides the state and configuration settings for that feature in lower layers. If you want a preprocessor or advanced setting to inherit its state and configuration from the base policy or a lower layer, set it to **Inherit**. Note that the Inherit selection is not available when you are working in the Settings or Advanced Settings page.

Color-coding on each layer summary page indicates as follows whether the effective configuration is in a higher, lower, or the current layer:

- red - the effective configuration is in a higher layer
- yellow - the effective configuration is in a lower layer
- unshaded - the effective configuration is in the current layer

Because the Settings and Advanced Settings pages are composite views of all relevant settings, these page do not use color coding to indicate the positions of effective configurations.

The system uses the configuration in the highest layer where the feature is enabled. Unless you explicitly modify the configuration, the system uses the default configuration. For example, if you enable and modify the network analysis DCE/RPC preprocessor in one layer, and you also enable but do not modify it in a higher layer, the system uses the default configuration in the higher layer.

The following table describes the actions available on the summary page for user-configurable layers.

Table 19-3 Layer Summary Page Actions

To...	You can...
modify the layer name or description	type a new value for **Name** or **Description**.
share the layer with other intrusion policies	select **Allow this layer to be used by other policies**. See Sharing Layers Between Policies, page 19-9 for more information.
enable or disable a preprocessor/advanced setting in the current layer	click **Enabled** or **Disabled** next to the feature. When you enable, a sublink to the configuration page appears beneath the layer name in the navigation panel, and an edit icon ( ) appears on the summary page next to the feature. Disabling removes the sublink and edit icon.
inherit the preprocessor/advanced setting state and configuration from the settings in the highest layer below the current layer	click **Inherit**. The page refreshes and, if the feature was enabled, the feature sublink in the navigation panel and the edit icon no longer appear.
access the configuration page for an enabled preprocessor/advanced setting	click the edit icon ( ) or the feature sublink to modify the current configuration. Note that the Back Orifice preprocessor has no user-configurable options.

To modify preprocessors/advanced settings in a user layer:

Step 1 While editing your policy, expand **Policy Layers** in the navigation panel, then click the name of the layer you want to modify.

The summary page for the layer appears.

Step 2 You can take any of the actions in the Layer Summary Page Actions table.

Step 3 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.



Customizing Traffic Preprocessing

Many of the advanced settings in an access control policy govern intrusion detection and prevention configurations that require specific expertise to configure. Advanced settings typically require little or no modification and are not common to every deployment.

This chapter explains how to set the following preferences:

- Setting the Default Intrusion Policy for Access Control, page 20-1 explains how to change the access control policy's default intrusion policy, which is used to initially inspect traffic before the system can determine exactly how to inspect that traffic
- Customizing Preprocessing with Network Analysis Policies, page 20-2 explains how to tailor certain traffic preprocessing options to specific security zones, and networks by assigning custom network analysis policies to preprocess matching traffic.

Other chapters describe policy-wide preprocessing and performance options for access control policies. For more information, see:

- Configuring Advanced Transport/Network Settings, page 24-1
- Tuning Preprocessing in Passive Deployments, page 25-1
- Tuning Intrusion Prevention Performance, page 11-6
- Tuning File and Malware Inspection Performance and Storage, page 11-16

Setting the Default Intrusion Policy for Access Control

License: Any

Each access control policy uses its *default intrusion policy* to initially inspect traffic before the system can determine exactly how to inspect that traffic. This is needed because sometimes the system must process the first few packets in a connection, **allowing them to pass**, before it can decide which access control rule (if any) will handle the traffic. However, so that these packets do not reach their destination uninspected, you can use an intrusion policy—called the default intrusion policy—to inspect them and generate intrusion events.

A default intrusion policy is especially useful when performing application control and URL filtering, because the system cannot identify applications or filter URLs before a connection is fully established between the client and the server. For example, if a packet matches all the other conditions in an access control rule with an application or URL condition, it and subsequent packets are allowed to pass until the connection is established and application or URL identification is complete, usually 3 to 5 packets.

The system inspects these allowed packets with the default intrusion policy, which can generate events and, if placed inline, block malicious traffic. After the system identifies the access control rule or default action that should handle the connection, the remaining packets in the connection are handled and inspected accordingly.

When you create an access control policy, its default intrusion policy depends on the default action you **first** chose. Initial default intrusion policies for access control are as follows:

- Balanced Security and Connectivity (a system-provided policy) is the default intrusion policy for an access control policy where you first chose the **Intrusion Prevention** default action.
- No Rules Active is the default intrusion policy for an access control policy where you first chose the **Block all traffic** default action. Although choosing this option disables intrusion inspection on the allowed packets described above, it can improve performance if you are not interested in intrusion data.

Note If you are not performing intrusion inspection, keep the No Rules Active policy as your default intrusion policy. For more information, see Troubleshooting Access Control Policies and Rules, page 4-12.

Note that if you change your default action after you create the access control policy, the default intrusion policy does **not** automatically change. To change it manually, use the access control policy's advanced options.

To change an access control policy's default intrusion policy:

Step 1 In the access control policy where you want to change the default intrusion policy, select the **Advanced** tab, then click the edit icon ( ) next to the Network Analysis and Intrusion Policies section.

The Network and Analysis Policies dialog box appears.

Step 2 From the **Intrusion Policy used before Access Control rule is determined** drop-down list, select a default intrusion policy. You can choose a system- or user-created policy.

Note that if you choose a user-created policy, you can click an edit icon ( ) to edit the policy in a new window. You cannot edit system-provided policies.

Caution Do **not** use `Experimental Policy 1` unless instructed to do so by a Cisco representative. Cisco uses this policy for testing.

Step 3 Click **OK** to save your changes.

You must apply the access control policy for your changes to take effect.

Customizing Preprocessing with Network Analysis Policies

License: Any

Network analysis policies govern how traffic is decoded and preprocessed so that it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt. This traffic preprocessing occurs after Security Intelligence blacklisting and traffic decryption, but before intrusion policies inspect packets in detail. By default, the system-provided Balanced Security and Connectivity network analysis policy applies to *all* traffic handled by an access control policy.

Tip The system-provided Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules.

A simple way to tune preprocessing is to create and use a custom network analysis policy as the default; see Creating a Custom Network Analysis Policy, page 21-2. Tuning options available vary by preprocessor.

For advanced users with complex deployments, you can create multiple network analysis policies, each tailored to preprocess traffic differently. Then, you can configure the system to use those policies to govern the preprocessing of traffic using different security zones or networks.

To accomplish this, you add custom *network analysis rules* to your access control policy. Each rule has:

- a set of rule conditions that identifies the specific traffic you want to preprocess
- an associated network analysis policy that you want to use to preprocess traffic that meets all the rules' conditions

When it is time for the system to preprocess traffic, it matches packets to network analysis rules in top-down order by rule number. Traffic that does not match any network analysis rules is preprocessed by the default network analysis policy.

Note If you disable a preprocessor but the system needs to evaluate preprocessed packets against an enabled intrusion or preprocessor rule, the system automatically enables and uses the preprocessor although it remains disabled in the network analysis policy interface. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task. Because preprocessing and intrusion inspection are so closely related, you **must** be careful that you allow the network analysis and intrusion policies examining a single packet to complement each other. For more information, see Limitations of Custom Policies, page 18-11.

For more information, see the following sections:

- Setting the Default Network Analysis Policy for Access Control, page 20-3
- Specifying Traffic to Preprocess Using Network Analysis Rules, page 20-4
- Managing Network Analysis Rules, page 20-7

Setting the Default Network Analysis Policy for Access Control

License: Any

By default, the system-provided Balanced Security and Connectivity network analysis policy applies to all traffic handled by an access control policy. If you add network analysis rules to tailor traffic preprocessing options, the default network analysis policy preprocesses all traffic not handled by those rules.

An access control policy's advanced settings allow you to change this default policy.

To change an access control policy's default network analysis policy:

Step 1 In the access control policy where you want to change the default network analysis policy, select the **Advanced** tab, then click the edit icon ( ) next to the Network Analysis and Intrusion Policies section.

The Network and Analysis Policies dialog box appears.

Step 2 From the **Default Network Analysis Policy** drop-down list, select a default network analysis policy. You can choose a system- or user-created policy.

Note that if you choose a user-created policy, you can click an edit icon ( ) to edit the policy in a new window. You cannot edit system-provided policies.

Caution Do **not** use `Experimental Policy 1` unless instructed to do so by a Cisco representative. Cisco uses this policy for testing.

Step 3 Click **OK** to save your changes,.

You must apply the access control policy for your changes to take effect.

Specifying Traffic to Preprocess Using Network Analysis Rules

License: Any

Within your access control policy's advanced settings, you can use network analysis rules to tailor preprocessing configurations to network traffic. Similar to access control rules, network analysis rules are numbered, starting at 1.

When it is time for the system to preprocess traffic, it matches packets to network analysis rules in top-down order by ascending rule number, and preprocesses traffic according to the first rule where all the rule's conditions match. The conditions you can add to a rule are described in the following table.

Table 20-1 Network Analysis Rule Condition Types

This Condition...	Matches Traffic...	Details
Zones	entering or leaving a device via an interface in a specific security zone	A security zone is a logical grouping of one or more interfaces according to your deployment and security policies. To build a zone condition, see Preprocessing Traffic Per Zone, page 20-5.
Networks	by its source or destination IP address	You can explicitly specify IP addresses. To build a network condition, see Preprocessing Traffic Per Network, page 20-6.

If you do not configure a particular condition for a rule, the system does not match traffic based on that criterion. For example, a rule with a network condition but no zone condition evaluates traffic based on its source or destination IP address, regardless of its ingress or egress interface. Traffic that does not match any network analysis rules is preprocessed by the default network analysis policy.

To add a custom network analysis rule:

Step 1 In the access control policy where you want to create custom preprocessing configurations, select the **Advanced** tab, then click the edit icon ( ) next to the Intrusion and Network Analysis Policies section.

The Network and Analysis Policies dialog box appears. If you have not added any custom network analysis rules, the module interface indicates that you have **No Custom Rules**, otherwise it displays how many you have configured.

Tip Click **Network Analysis Policy List** to display the Network Analysis Policy page in a new window. Use this page to view and edit your custom network analysis policies; see Managing Network Analysis Policies, page 21-3

Step 2 Next to **Network Analysis Rules**, click the statement that indicates how many custom rules you have.

The dialog box expands to show the custom rules, if any.

Step 3 Click **Add Rule**.

The network analysis rule editor appears.

Step 4 Build your rule's conditions. You can restrict NAP preprocessing using the following criteria:

- Preprocessing Traffic Per Zone, page 20-5
- Preprocessing Traffic Per Network, page 20-6

Step 5 Associate a network analysis policy with the rule by clicking the **Network Analysis** tab and choosing a policy from the **Network Analysis Policy** drop-down list.

The system uses the network analysis policy you choose to preprocess traffic that meets all the rule's conditions. Note that if you choose a user-created policy, you can click an edit icon ( ) to edit the policy in a new window. You cannot edit system-provided policies.

Caution Do **not** use `Experimental Policy 1` unless instructed to do so by a Cisco representative. Cisco uses this policy for testing.

Step 6 Click **Add**.

The rule is added after any other rules. To change the rule's evaluation order, see Managing Network Analysis Rules, page 20-7.

Preprocessing Traffic Per Zone

License: Any

Zone conditions in network analysis rules allow you to preprocess traffic by its source and destination security zones. A security zone is a grouping of one or more interfaces. For more information on creating zones, see Working with Security Zones, page 2-32.

You can add a maximum of 50 zones to each of the **Source Zones** and **Destination Zones** in a single zone condition:

- To match traffic *leaving* the device from an interface in the zone, add that zone to **Destination Zones**. Note that because devices deployed passively do not transmit traffic, you cannot use a zone comprised of passive interfaces in a **Destination Zone** condition.
- To match traffic *entering* the device from an interface in the zone, add that zone to **Source Zones**.

If you add both source and destination zone conditions to a rule, matching traffic must originate from one of the specified source zones and egress through one of the destination zones.

Warning icons ( ) indicate invalid configurations, such as zones that contain no interfaces. For details, see Troubleshooting Access Control Policies and Rules, page 4-12.

To preprocess traffic by zone:

Step 1 In the access control policy where you want to preprocess traffic by zone, create a new network analysis rule or edit an existing rule.

For detailed instructions, see Specifying Traffic to Preprocess Using Network Analysis Rules, page 20-4.

Step 2 In the network analysis rule editor, select the **Zones** tab.

The Zones tab appears.

Step 3 Find and select the zones you want to add from the **Available Zones**.

To search for zones to add, click the **Search by name** prompt above the **Available Zones** list, then type a zone name. The list updates as you type to display matching zones.

Click to select a zone. To select multiple zones, use the **Shift** and **Ctrl** keys, or right-click and then select **Select All**.

Step 4 Click **Add to Source** or **Add to Destination** to add the selected zones to the appropriate list.

You can also drag and drop selected zones.

Step 5 Save or continue editing the rule.

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Preprocessing Traffic Per Network

License: Any

Network conditions in network analysis rules allow you to preprocess traffic by its source and destination IP address. You can manually specify the source and destination IP addresses for the traffic you want to preprocess, or you can configure network conditions with network objects, which are reusable and associate a name with one or more IP addresses and address blocks.

Tip After you create a network object, you can use it not only to build network analysis rules, but also to represent IP addresses in various other places in the system's module interface. You can create these objects using the object manager; you can also create network objects on-the-fly while you are configuring network analysis rules. For more information, see Working with Network Objects, page 2-3.

You can add a maximum of 50 items to each of the **Source Networks** and **Destination Networks** in a single network condition:

- To match traffic from an IP address, configure **Source Networks**.
- To match traffic to an IP address, configure **Destination Networks**.

If you add both source and destination network conditions to a rule, matching traffic must originate from one of the specified IP addresses and be destined for one of the destination IP addresses.

When building a network condition, warning icons ( ) indicate invalid configurations. For details, see Troubleshooting Access Control Policies and Rules, page 4-12.

To preprocess traffic by network:

Step 1 In the access control policy where you want to preprocess traffic by network, create a new network analysis rule or edit an existing rule.

For detailed instructions, see Specifying Traffic to Preprocess Using Network Analysis Rules, page 20-4.

Step 2 In the network analysis rule editor, select the **Networks** tab.

The Networks tab appears.

Step 3 Find and select the networks you want to add from the **Available Networks**, as follows:

- To add a network object on the fly, which you can then add to the condition, click the add icon ( ) above the **Available Networks** list; see Working with Network Objects, page 2-3.
- To search for networks to add, click the **Search by name or value** prompt above the **Available Networks** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.

To select an object, click it. To select multiple objects, use the Shift and Ctrl keys, or right-click and then select **Select All**.

Step 4 Click **Add to Source** or **Add to Destination** to add the selected objects to the appropriate list.

You can also drag and drop selected objects.

Step 5 Add any source or destination IP addresses or address blocks that you want to specify manually.

Click the **Enter an IP address** prompt below the **Source Networks** or **Destination Networks** list; then type an IP address or address block and click **Add**.

Step 6 Save or continue editing the rule.

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Managing Network Analysis Rules

License: Any

A network analysis rule is simply a set of configurations and conditions that specifies how you preprocess traffic that matches those qualifications. You create and edit network analysis rules in the advanced options in an existing access control policy. Each rule belongs to only one policy.

To edit a custom network analysis rule:

Step 1 In the access control policy where you want to change your custom preprocessing configurations, select the **Advanced** tab, then click the edit icon ( ) next to the Intrusion and Network Analysis Policies section.

The Network and Analysis Policies dialog box appears. If you have not added any custom network analysis rules, the module interface indicates that you have **No Custom Rules**; otherwise, it displays how many you have configured.

Step 2 Next to **Network Analysis Rules**, click the statement that indicates how many custom rules you have.

The dialog box expands to show the custom rules, if any.

Step 3 Edit your custom rules. You have the following options:

- To edit a rule's conditions, or change the network analysis policy invoked by the rule, click the edit icon ( ) next to the rule.
- To change a rule's order of evaluation, click and drag the rule to the correct location. To select multiple rules, use the Shift and Ctrl keys.
- To delete a rule, click the delete icon ( ) next to the rule.

Step 4 Click **OK** to save your changes.

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.



Getting Started with Network Analysis Policies

Network analysis policies govern many traffic preprocessing options, and are invoked by advanced settings in your access control policy. Network analysis-related preprocessing occurs after Security Intelligence blacklisting and SSL decryption, but before access control rules inspect packets in detail, and before any intrusion or file inspection begins.

By default, the system uses the *Balanced Security and Connectivity* network analysis policy to preprocess all traffic handled by an access control policy. However, you can choose a different default network analysis policy to perform this preprocessing. For your convenience, the system provides a choice of several non-modifiable network analysis policies, which are tuned for a specific balance of security and connectivity by the Vulnerability Research Team (VRT). You can also replace this default policy with a custom network analysis policy with custom preprocessing settings.

Tip System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules. Understanding Network Analysis and Intrusion Policies, page 18-1 provides an overview of how network analysis and intrusion policies work together to examine your traffic, as well as some basics on using the navigation panel, resolving conflicts, and committing changes.

You can also tailor traffic preprocessing options to specific security zones, and networks by creating multiple custom network analysis policies, then assigning them to preprocess different traffic.

Note Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task. Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. The system does **not** coordinate the policies for you, and uses default options in cases of misconfiguration. For more information, see Limitations of Custom Policies, page 18-11.

This chapter explains how to create a simple custom network analysis policy. This chapter also contains basic information on managing network analysis policies: editing, comparing, and so on. For more information, see:

- Creating a Custom Network Analysis Policy, page 21-2
- Managing Network Analysis Policies, page 21-3

- Allowing Preprocessors to Affect Traffic in Inline Deployments, page 21-5
- Generating a Report of Current Network Analysis Settings, page 21-8
- Comparing Two Network Analysis Policies or Revisions, page 21-9

Creating a Custom Network Analysis Policy

License: Any

When you create a new network analysis policy you must give it a unique name, specify a base policy, and choose an *inline mode*.

The base policy defines the network analysis policy's default settings. Modifying a setting in the new policy overrides—but does not change—the settings in the base policy. You can use either a system-provided or custom policy as your base policy. For more information, see Understanding the Base Layer, page 19-2.

The network analysis policy's inline mode allows preprocessors to modify (normalize) and drop traffic to minimize the chances of attackers evading detection. Note that in passive deployments, the system cannot affect traffic flow regardless of the inline mode. For more information, see Allowing Preprocessors to Affect Traffic in Inline Deployments, page 21-5.

To create a network analysis policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click **Create Policy**.

If you have unsaved changes in another policy, click **Cancel** when prompted to return to the Network Analysis Policy page. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Create Network Analysis Policy pop-up window appears.

Step 7 Give the policy a unique **Name** and, optionally, a **Description**.

Step 8 Specify the initial **Base Policy**.

You can either use either a system-provided or custom policy as your base policy.

Caution Do **not** use `Experimental Policy 1` unless instructed to do so by a Cisco representative. Cisco uses this policy for testing.

Step 9 Specify whether you want to allow preprocessors to affect traffic in an inline deployment:

- To allow preprocessors to affect traffic, enable **Inline Mode**.
- To prevent preprocessors from affecting traffic, disable **Inline Mode**.

Step 10 Create the policy:

- Click **Create Policy** to create the new policy and return to the Network Analysis Policy page. The new policy has the same settings as its base policy.
- Click **Create and Edit Policy** to create the policy and open it for editing in the advanced network analysis policy editor; see Editing Network Analysis Policies, page 21-3.

Managing Network Analysis Policies

License: Any

On the Network Analysis Policy page you can view your current custom network analysis policies, along with the following information:

- the time and date the policy was last modified (in local time) and the user who modified it
- whether the **Inline Mode** setting is enabled, which allows preprocessors to affect traffic
- which access control policies are using the network analysis policy to preprocess traffic
- whether a policy has unsaved changes, as well as information about who (if anyone) is currently editing the policy

Options on the Network Analysis Policy page allow you to take the actions in the following table.

Table 21-1 Network Analysis Policy Management Actions

To...	You can...	See...
create a new network analysis policy	click **Create Policy**.	Creating a Custom Network Analysis Policy, page 21-2.
edit an existing network analysis policy	click the edit icon ( ).	Editing Network Analysis Policies, page 21-3.
view a PDF report that lists the current configuration settings in a network analysis policy	click the report icon ( ).	Generating a Report of Current Network Analysis Settings, page 21-8
compare the settings of two network analysis policies or two revisions of the same policy	click **Compare Policies**.	Comparing Two Network Analysis Policies or Revisions, page 21-9.
delete a network analysis policy	click the delete icon ( ), then confirm that you want to delete the policy. You cannot delete a network analysis policy if an access control policy references it.	

Editing Network Analysis Policies

License: Any

When you create a new network analysis policy, it has the same settings as its base policy. The following table lists the most common actions you can take to tailor the new policy to your needs:

Table 21-2 Network Analysis Policy Editing Actions

To...	You can...	See...
allow preprocessors to modify or drop traffic	select the **Inline Mode** check box on the Policy Information page.	Allowing Preprocessors to Affect Traffic in Inline Deployments, page 21-5
change the base policy	select a base policy from the **Base Policy** drop-down list on the Policy Information page.	Changing the Base Policy, page 19-3
view the settings in the base policy	click **Manage Base Policy** on the Policy Information page.	Understanding the Base Layer, page 19-2
enable, disable, or edit the settings for a preprocessor	click **Settings** in the navigation panel.	Configuring Preprocessors in a Network Analysis Policy, page 21-6
manage policy layers	click **Policy Layers** in the navigation panel.	Using Layers in a Network Analysis or Intrusion Policy, page 19-1

When tailoring a network analysis policy, especially when disabling preprocessors, keep in mind that some preprocessors and intrusion rules require that traffic first be decoded or preprocessed in a certain way. If you disable a required preprocessor, the system automatically uses it with its current settings, although the preprocessor remains disabled in the network analysis policy module interface.

Note Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task. For more information, see Limitations of Custom Policies, page 18-11.

The system caches one network analysis policy per user. While editing a network analysis policy, if you select any menu or other path to another page, your changes stay in the system cache even if you leave the page. In addition to the actions you can perform in the table above, Understanding Network Analysis and Intrusion Policies, page 18-1 provides information on using the navigation panel, resolving conflicts, and committing changes.

To edit a network analysis policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the network analysis policy you want to configure.

The network analysis policy editor appears, focused on the Policy Information page and with a navigation panel on the left.

Step 7 Edit your policy. Take any of the actions summarized above.

Step 8 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Allowing Preprocessors to Affect Traffic in Inline Deployments

License: Any

In an inline deployment, some preprocessors can modify and block traffic. For example:

- The inline normalization preprocessor normalizes packets to prepare them for analysis by other preprocessors and the intrusion rules engine. You can also use the preprocessor's **Block Unrecoverable TCP Header Anomalies** and **Allow These TCP Options** options to block certain packets. For more information, see Normalizing Inline Traffic, page 24-6.
- The system can drop packets with invalid checksums; see Verifying Checksums, page 24-4.
- The system can drop packets matching rate-based attack prevention settings; see Preventing Rate-Based Attacks, page 28-9.

For a preprocessor configured in the network analysis policy to affect traffic, you must also enable and correctly configure the preprocessor, as well as correctly deploy the device inline. Finally, you must enable the network analysis policy's **Inline Mode** setting.

If you want to assess how your configuration would function in an inline deployment without actually modifying traffic, you can disable inline mode. Note that in passive deployments, the system cannot affect traffic regardless of the inline mode.

Tip In an inline deployment, Cisco recommends that you enable inline mode and configure the inline normalization preprocessor with the **Normalize TCP Payload** option enabled. In a passive deployment,Cisco recommends you configure adaptive profiles.

To allow preprocessors to affect traffic in an inline deployment:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

The Policy Information page appears.

Step 7 Specify whether you want to allow preprocessors to affect traffic:

- To allow preprocessors to affect traffic, enable **Inline Mode**.
- To prevent preprocessors from affecting traffic, disable **Inline Mode**.

Step 8 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Configuring Preprocessors in a Network Analysis Policy

License: Any

Preprocessors prepare traffic to be further inspected by normalizing traffic and identifying protocol anomalies. Preprocessors generate preprocessor event when packets trigger preprocessor options that you configure. The base policy for your network analysis policy determines which preprocessors are enabled by default and the default configuration for each.

When you select **Settings** in the navigation panel of a network analysis policy, the policy lists its preprocessors by type. On the Settings page, you can enable or disable preprocessors in your network analysis policy, as well as access preprocessor configuration pages.

A preprocessor must be enabled for you to configure it. When you enable a preprocessor, a sublink to the configuration page for the preprocessor appears beneath the **Settings** link in the navigation panel, and an **Edit** link to the configuration page appears next to the preprocessor on the Settings page.

Tip To revert a preprocessor's configuration to the settings in the base policy, click **Revert to Defaults** on a preprocessor configuration page. When prompted, confirm that you want to revert.

When you disable a preprocessor, the sublink and **Edit** link no longer appear, but your configurations are retained. Note that to perform their particular analysis, many preprocessors and intrusion rules require that traffic first be decoded or preprocessed in a certain way. If you disable a required preprocessor, the system automatically uses it with its current settings, although the preprocessor remains disabled in the network analysis policy module interface.

Note In most cases, preprocessors require specific expertise to configure and typically require little or no modification. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task. Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. For more information, see Limitations of Custom Policies, page 18-11.

Modifying a preprocessor configuration requires an understanding of the configuration and its potential impact on your network. The following sections provide links to specific configuration details for each preprocessor.

Application Layer Preprocessors

Application-layer protocol decoders normalize specific types of packet data into formats that the intrusion rules engine can analyze.

Table 21-3 Application Layer Preprocessor Settings

For information on...	See...
DCE/RPC Configuration	Decoding DCE/RPC Traffic, page 22-2
DNS Configuration	Detecting Exploits in DNS Name Server Responses, page 22-14
FTP and Telnet Configuration	Decoding FTP and Telnet Traffic, page 22-18
HTTP Configuration	Decoding HTTP Traffic, page 22-31
Sun RPC Configuration	Using the Sun RPC Preprocessor, page 22-46
SIP Configuration	Decoding the Session Initiation Protocol, page 22-48
GTP Command Channel Configuration	Configuring the GTP Command Channel, page 22-52
IMAP Configuration	Decoding IMAP Traffic, page 22-54
POP Configuration	Decoding POP Traffic, page 22-57
SMTP Configuration	Decoding SMTP Traffic, page 22-60
SSH Configuration	Detecting Exploits Using the SSH Preprocessor, page 22-67
SSL Configuration	Using the SSL Preprocessor, page 22-71

SCADA Preprocessors

The Modbus and DNP3 preprocessors detect traffic anomalies and provide data to the intrusion rules engine for inspection.

Table 21-4 SCADA Preprocessor Settings

For information on...	See...
Modbus Configuration	Configuring the Modbus Preprocessor, page 23-1
DNP3 Configuration	Configuring the DNP3 Preprocessor, page 23-3

Transport/Network Layer Preprocessors

Network and transport layers preprocessors detect exploits at the network and transport layers. Before packets are sent to preprocessors, the packet decoder converts packet headers and payloads into a format that can be easily used by the preprocessors and the intrusion rules engine; it also detects various anomalous behaviors in packet headers.

Table 21-5 Transport and Network Layer Preprocessor Settings

For information on...	See...
Checksum Verification	Verifying Checksums, page 24-4
Inline Normalization	Normalizing Inline Traffic, page 24-6
IP Defragmentation	Defragmenting IP Packets, page 24-11
Packet Decoding	Understanding Packet Decoding, page 24-16

Table 21-5 Transport and Network Layer Preprocessor Settings (continued)

For information on...	See...
TCP Stream Configuration	Using TCP Stream Preprocessing, page 24-20
UDP Stream Configuration	Using UDP Stream Preprocessing, page 24-31

Note that some advanced transport and network preprocessor settings apply globally to all networks and zones where you apply your access control policy. You configure these advanced settings in an access control policy rather than in a network analysis policy; see Configuring Advanced Transport/Network Settings, page 24-1.

Specific Threat Detection

The Back Orifice preprocessor analyzes UDP traffic for the Back Orifice magic cookie. The portscan detector can be configured to report scan activity. Rate-based attack prevention can help you protect your network against SYN floods and an extreme number of simultaneous connections designed to overwhelm your network.

Table 21-6 Specific Threat Detection Settings

For information on...	See...
Back Orifice Detection	Detecting Back Orifice, page 28-1
Portscan Detection	Detecting Portscans, page 28-3
Rate-Based Attack Prevention	Preventing Rate-Based Attacks, page 28-9

Note that you configure the sensitive data preprocessor, which detects sensitive data such as credit card numbers and Social Security numbers in ASCII text, in intrusion policies. For more information, see Detecting Sensitive Data, page 28-19.

Generating a Report of Current Network Analysis Settings

License: Any

A network analysis policy report is a record of the policy configuration at a specific point in time. The system combines the settings in the base policy with the settings of the policy layers, and makes no distinction between which settings originated in the base policy or policy layer.

You can use the report, which contains the following information, for auditing purposes or to inspect the current configuration.

Table 21-7 Network Analysis Policy Report Sections

Section	Description
Policy Information	Provides the name and description of the policy, the name of the user who last modified the policy, and the date and time the policy was last modified. Also indicates whether inline normalization can be enabled, the current rule update version, and whether the base policy is locked to the current rule update.
Settings	Lists all enabled preprocessor settings and their configurations.

You can also generate a comparison report that compares two network analysis policies, or two revisions of the same policy. For more information, see Comparing Two Network Analysis Policies or Revisions, page 21-9.

To view a network analysis policy report:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the report icon ( ) next to the policy for which you want to generate a report. Remember to commit any changes before you generate a network analysis policy report; only committed changes appear in the report.

The system generates the report. You are prompted to save the report to your computer.

Comparing Two Network Analysis Policies or Revisions

License: Any

To review policy changes for compliance with your organization's standards or to optimize system performance, you can examine the differences between two network analysis policies. You can compare any two network analysis policies or two revisions of the same network analysis policy. Optionally, after you compare, you can then generate a PDF report to record the differences between the two policies or policy revisions.

There are two tools you can use to compare network analysis policies or policy revisions:

- The comparison view displays only the differences between two network analysis policies or network analysis policy revisions in a side-by-side format; the name of each policy or policy revision appears in the title bar on the left and right sides of the comparison view.

  You can use this to view and navigate both policy revisions on the module interface, with their differences highlighted.

- The comparison report creates a record of only the differences between two network analysis policies or network analysis policy revisions in a format similar to the network analysis policy report, but in PDF format.

  You can use this to save, copy, print and share your policy comparisons for further examination.

For more information on understanding and using the policy comparison tools, see:

- Using the Network Analysis Policy Comparison View, page 21-10

- Using the Network Analysis Policy Comparison Report, page 21-10

Using the Network Analysis Policy Comparison View

License: Any

The comparison view displays both policies or policy revisions in a side-by-side format, with each policy or policy revision identified by name in the title bar on the left and right sides of the comparison view. The time of last modification and the last user to modify are displayed with the policy name.

Differences between the two policies are highlighted:

- Blue indicates that the highlighted setting is different in the two policies, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one policy but not the other.

You can perform any of the actions in the following table.

Table 21-8 Network Analysis Policy Comparison View Actions

To...	You can...
navigate individually through changes	click **Previous** or **Next** above the title bar. The double-arrow icon ( ) centered between the left and right sides moves, and the **Difference** number adjusts to identify which difference you are viewing.
generate a new policy comparison view	click **New Comparison.** The **Select Comparison** window appears. See Using the Network Analysis Policy Comparison Report, page 21-10 for more information.
generate a policy comparison report	click **Comparison Report**. The policy comparison report creates a PDF document that lists only the differences between the two policies or policy revisions.

Using the Network Analysis Policy Comparison Report

License: Any

A network analysis policy comparison report is a record of all differences between two network analysis policies or two revisions of the same network analysis policy identified by the network analysis policy comparison view, presented as a PDF. You can use this report to further examine the differences between two network analysis policy configurations and to save and disseminate your findings.

You can generate a network analysis policy comparison report from the comparison view for any policies to which you have access. Remember to save any changes before you generate a policy report; only saved changes appear in the report.

The format of the policy comparison report is the same as the policy report with one exception: the policy report contains all configurations in the policy, and the policy comparison report lists only those configurations that differ between the policies. A network analysis policy comparison report contains the sections described in Table 21-7 on page 21-8.

Tip You can use a similar procedure to compare SSL, access control, intrusion, or file policies.

To compare two network analysis policies or policy revisions:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click **Compare Policies**.

The Select Comparison window appears.

Step 7 From the **Compare Against** drop-down list, select the type of comparison you want to make:

- To compare two different policies, select **Other Policy**.

  The page refreshes and the Policy A and Policy B drop-down lists appear.
- To compare two revisions of the same policy, select **Other Revision.**

  The page refreshes and the Policy, Revision A, and Revision B drop-down lists appear.

Step 8 Depending on the comparison type you selected, you have the following choices:

- If you are comparing two different policies, select the policies you want to compare from the Policy A and Policy B drop-down lists.
- If you are comparing two revision of the same policy, select the Policy, then select the timestamped revisions you want to compare from the Revision A and Revision B drop-down lists.

Step 9 Click **OK** to display the policy comparison view.

The comparison view appears.

Step 10 Optionally, click **Comparison Report** to generate the network analysis policy comparison report.

The network analysis policy comparison report appears. You are prompted to save the report to your computer.



Using Application Layer Preprocessors

You configure application layer preprocessors in a network analysis policy, which prepares traffic for inspection using the rules enabled in an intrusion policy. See Understanding Network Analysis and Intrusion Policies, page 18-1 for more information.

Application-layer protocols can represent the same data in a variety of ways. Cisco provides application layer protocol decoders that normalize specific types of packet data into formats that the intrusion rules engine can analyze. Normalizing application-layer protocol encodings allows the rules engine to effectively apply the same content-related rules to packets whose data is represented differently and obtain meaningful results.

Note that preprocessors do not generate events in most cases unless you enable the accompanying preprocessor rules in an intrusion policy. See Setting Rule States, page 27-19 for more information.

See the following sections for more information:

- Decoding DCE/RPC Traffic, page 22-2 describes the DCE/RPC preprocessor and explains how to configure it to prevent evasion attempts and detect anomalies in DCE/RPC traffic.
- Detecting Exploits in DNS Name Server Responses, page 22-14 describes the DNS preprocessor and explains how to configure it to detect any of three specific exploits in DNS name server responses.
- Decoding FTP and Telnet Traffic, page 22-18 describes the FTP/Telnet decoder and explains how to configure it to normalize and decode FTP and Telnet traffic.
- Decoding HTTP Traffic, page 22-31 describes the HTTP decoder and explains how to configure it to normalize HTTP traffic.
- Using the Sun RPC Preprocessor, page 22-46 describes the RPC decoder and explains how to configure it to normalize RPC traffic.
- Decoding the Session Initiation Protocol, page 22-48 explains how you can use the SIP preprocessor to decode and detect anomalies in SIP traffic.
- Configuring the GTP Command Channel, page 22-52 explains how you can use the GTP preprocessor to provide the rules engine with GTP command channel messages extracted by the packet decoder.
- Decoding IMAP Traffic, page 22-54 explains how you can use the IMAP preprocessor to decode and detect anomalies in IMAP traffic.
- Decoding POP Traffic, page 22-57 explains how you can use the POP preprocessor to decode and detect anomalies in POP traffic.
- Decoding SMTP Traffic, page 22-60 describes the SMTP decoder and explains how to configure it to decode and normalize SMTP traffic.

- Detecting Exploits Using the SSH Preprocessor, page 22-67 explains how to identify and process exploits in SSH-encrypted traffic.
- Using the SSL Preprocessor, page 22-71 explains how you can use the SSL preprocessor to identify encrypted traffic and eliminate false positives by stopping inspection of that traffic.
- Configuring SCADA Preprocessing, page 23-1 explains how you can use the Modbus and DNP3 preprocessors to detect anomalies in corresponding traffic and provide data to the intrusion rules engine for inspection of certain protocol fields.

Decoding DCE/RPC Traffic

License: Protection

The DCE/RPC protocol allows processes on separate network hosts to communicate as if the processes were on the same host. These inter-process communications are commonly transported between hosts over TCP and UDP. Within the TCP transport, DCE/RPC might also be further encapsulated in the Windows Server Message Block (SMB) protocol or in Samba, an open-source SMB implementation used for inter-process communication in a mixed environment comprised of Windows and UNIX- or Linux-like operating systems. In addition, Windows IIS web servers on your network might use IIS RPC over HTTP, which provides distributed communication through a firewall, to proxy TCP-transported DCE/RPC traffic.

Note that descriptions of DCE/RPC preprocessor options and functionality include the Microsoft implementation of DCE/RPC known as MSRPC; descriptions of SMB options and functionality refer to both SMB and Samba.

Although most DCE/RPC exploits occur in DCE/RPC client requests targeted for DCE/RPC servers, which could be practically any host on your network that is running Windows or Samba, exploits can also occur in server responses. The DCE/RPC preprocessor detects DCE/RPC requests and responses encapsulated in TCP, UDP, and SMB transports, including TCP-transported DCE/RPC using version 1 RPC over HTTP. The preprocessor analyzes DCE/RPC data streams and detects anomalous behavior and evasion techniques in DCE/RPC traffic. It also analyzes SMB data streams and detects anomalous SMB behavior and evasion techniques.

The DCE/RPC preprocessor also desegments SMB and defragments DCE/RPC in addition to the IP defragmentation provided by the IP defragmentation preprocessor and the TCP stream reassembly provided by the TCP stream preprocessor. See Using TCP Stream Preprocessing, page 24-20 and Defragmenting IP Packets, page 24-11.

Finally, the DCE/RPC preprocessor normalizes DCE/RPC traffic for processing by the rules engine. See DCE/RPC Keywords, page 30-59 for information on using specific DCE/RPC rule keywords to detect DCE/RPC services, operations, and stub data.

You configure the DCE/RPC preprocessor by modifying any of the global options that control how the preprocessor functions, and by specifying one or more target-based server policies that identify the DCE/RPC servers on your network by IP address and by either the Windows or Samba version running on them:

You must enable DCE/RPC preprocessor rules, which have a generator ID (GID) of 132 or 133, if you want these rules to generate events. See Setting Rule States, page 27-19 for more information.

See the following sections for more information:

- Selecting Global DCE/RPC Options, page 22-3
- Understanding Target-Based DCE/RPC Server Policies, page 22-4
- Understanding DCE/RPC Transports, page 22-5

- Selecting DCE/RPC Target-Based Policy Options, page 22-8
- Configuring the DCE/RPC Preprocessor, page 22-11

Selecting Global DCE/RPC Options

License: Protection

Global DCE/RPC preprocessor options control how the preprocessor functions. Except for the **Memory Cap Reached** option, modifying these options could have a negative impact on performance or detection capability. You should not modify them unless you have a thorough understanding of the preprocessor and the interaction between the preprocessor and enabled DCE/RPC rules. In particular, make sure that the **Maximum Fragment Size** option and **Reassembly Threshold** option are greater than or equal to the depth to which the rules need to detect. For more information, see Constraining Content Matches, page 30-17 and Using Byte_Jump and Byte_Test, page 30-30.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Maximum Fragment Size

When **Enable Defragmentation** is selected, specifies the maximum DCE/RPC fragment length allowed from 1514 to 65535 bytes. The preprocessor truncates larger fragments for processing purposes to the specified size before defragmenting but does not alter the actual packet. A blank field disables this option.

Reassembly Threshold

When **Enable Defragmentation** is selected, 0 disables this option, or 1 to 65535 bytes specifies a minimum number of fragmented DCE/RPC bytes and, if applicable, segmented SMB bytes to queue before sending a reassembled packet to the rules engine. A low value increases the likelihood of early detection but could have a negative impact on performance. You should test for performance impact if you enable this option.

Enable Defragmentation

Specifies whether to defragment fragmented DCE/RPC traffic. When disabled, the preprocessor still detects anomalies and sends DCE/RPC data to the rules engine, but at the risk of missing exploits in fragmented DCE/RPC data.

Although this option provides the flexibility of not defragmenting DCE/RPC traffic, most DCE/RPC exploits attempt to take advantage of fragmentation to hide the exploit. Disabling this option would bypass most known exploits, resulting in a large number of false negatives.

Memory Cap Reached

Detects when the maximum memory limit allocated to the preprocessor is reached or exceeded. When the maximum memory cap is reached or exceeded, the preprocessor frees all pending data associated with the session that caused the memory cap event and ignores the rest of that session.

You can enable rule 133:1 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Auto-Detect Policy on SMB Session

Detects the Windows or Samba version that is identified in SMB `Session Setup AndX` requests and responses. When the detected version is different from the Windows or Samba version configured for the **Policy** configuration option, the detected version overrides the configured version for that

session only. See Understanding Target-Based DCE/RPC Server Policies, page 22-4 for more information.

For example, if you set **Policy** to Windows XP and the preprocessor detects Windows Vista, the preprocessor uses a Windows Vista policy for that session. Other settings remain in effect.

When the DCE/RPC transport is not SMB (that is, when the transport is TCP or UDP), the version cannot be detected and the policy cannot be automatically configured.

To enable this option, select one of the following from the drop-down list:

- Select **Client** to inspect server-to-client traffic for the policy type.
- Select **Server** to inspect client-to-server traffic for the policy type.
- Select **Both** to inspect server-to-client and client-to-server traffic for the policy type.

Understanding Target-Based DCE/RPC Server Policies

License: Protection

You can create one or more target-based server policies to configure the DCE/RPC preprocessor to inspect DCE/RPC traffic the same as a specified type of server would process it. Target-based policy configuration includes identifying the Windows or Samba version running on hosts you identify on your network, enabling transport protocols and specifying the ports carrying DCE/RPC traffic to those hosts, and setting other server-specific options.

Windows and Samba DCE/RPC implementations differ significantly. For example, all versions of Windows use the DCE/RPC context ID in the first fragment when defragmenting DCE/RPC traffic, and all versions of Samba use the context ID in the last fragment. As another example, Windows Vista uses the opnum (operation number) header field in the first fragment to identify a specific function call, and Samba and all other Windows versions use the opnum field in the last fragment.

There are also significant differences in Windows and Samba SMB implementations. For example, Windows recognizes the SMB OPEN and READ commands when working with named pipes, but Samba does not recognize these commands.

When you enable the DCE/RPC preprocessor, you automatically enable a default target-based policy. Optionally, you can add target-based policies that target other hosts running different Windows or Samba versions by selecting the correct version from the **Policy** drop-down list. The default target-based policy applies to any host not included in another target-based policy.

In each target-based policy, you can enable one or more transports and specify *detection ports* for each. You can also enable and specify *auto-detection ports*. See Understanding DCE/RPC Transports, page 22-5 for more information.

You can also configure other target-based policy options. You can set the preprocessor to detect when there is an attempt to connect to one or more shared SMB resources that you identify. You can configure the preprocessor to detect files in SMB traffic, and to inspect a specified number of bytes in a detected file. You can also modify an advanced option that should be modified only by a user with SMB protocol expertise; this option lets you set the preprocessor to detect when a number of chained SMB AndX commands exceed a specified maximum number.

In each target-based policy, you can:

- enable one or more transports and specify *detection ports* for each.
- enable and specify *auto-detection ports*. See Understanding DCE/RPC Transports, page 22-5 for more information.

- set the preprocessor to detect when there is an attempt to connect to one or more shared SMB resources that you identify.
- configure the preprocessor to detect files in SMB traffic, and to inspect a specified number of bytes in a detected file.
- modify an advanced option that should be modified only by a user with SMB protocol expertise; this option lets you set the preprocessor to detect when a number of chained SMB AndX commands exceed a specified maximum number.

Note that you can enable the **Auto-Detect Policy on SMB Session** global option to automatically override the policy type configured for a targeted policy on a per session basis when SMB is the DCE/RPC transport. See Auto-Detect Policy on SMB Session, page 22-3.

In addition to enabling SMB traffic file detection in the DCE/RPC preprocessor, you can configure a file policy to optionally capture and block these files. Within that policy, you must create a file rule with an **Action** of **Detect Files** or **Block Files** and a selected **Application Protocol** of **Any** or **NetBIOS-ssn (SMB)**. See Creating a File Policy, page 35-9 and Working with File Rules, page 35-10 for more information.

Understanding DCE/RPC Transports

License: Protection

In each target-based policy, you can enable one or more of the TCP, UDP, SMB, and RPC over HTTP transports. When you enable a transport, you must also specify one or more *detection ports*, that is, ports that are known to carry DCE/RPC traffic. Optionally, you can also enable and specify *auto-detection ports*, that is, ports that the preprocessor tests first to determine if they carry DCE/RPC traffic and continues processing only when it detects DCE/RPC traffic.

Cisco recommends that you use the default detection ports, which are either well-known ports or otherwise commonly-used ports for each protocol. You would add detection ports only if you detected DCE/RPC traffic on a non-default port.

When you enable auto-detection ports, ensure that they are set to the port range from 1024 to 65535 to cover the entire ephemeral port range. Note that it is unlikely that you would enable or specify auto-detection ports for the RPC over HTTP Proxy Auto-Detect Ports option or the SMB Auto-Detect Ports option because there is little likelihood that traffic for either would occur or even be possible except on the specified default detection ports. Note also that auto-detection occurs only for ports not already identified by transport detection ports. See Selecting DCE/RPC Target-Based Policy Options, page 22-8 for recommendations for enabling or disabling auto-detection ports for each transport.

You can specify ports for one or more transports in any combination in a Windows target-based policy to match the traffic on your network, but you can only specify ports for the SMB transport in a Samba target-based policy.

Note that you must enable at least one DCE/RPC transport in the default target-based policy except when you have added a DCE/RPC target-based policy that has at least one transport enabled. For example, you might want to specify the hosts for all DCE/RPC implementations and not have the default target-based policy apply to unspecified hosts, in which case you would not enable a transport for the default target-based policy.

See the following sections for more information:

- Understanding Connectionless and Connection-Oriented DCE/RPC Traffic, page 22-6
- Understanding the RPC over HTTP Transport, page 22-7

Understanding Connectionless and Connection-Oriented DCE/RPC Traffic

License: Protection

DCE/RPC messages comply with one of two distinct DCE/RPC Protocol Data Unit (PDU) protocols:

- the connection-oriented DCE/RPC PDU protocol

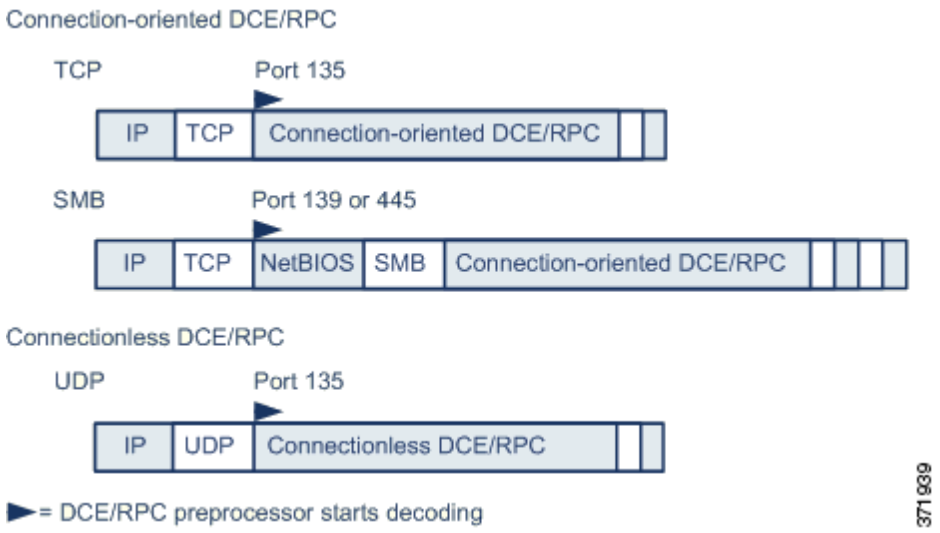
  The DCE/RPC preprocessor detects connection-oriented DCE/RPC in the TCP, SMB, and RPC over HTTP transports.

- the connectionless DCE/RPC PDU protocol

  The DCE/RPC preprocessor detects connectionless DCE/RPC in the UDP transport.

The two DCE/RPC PDU protocols have their own unique headers and data characteristics. For example, the connection-oriented DCE/RPC header length is typically 24 bytes and the connectionless DCE/RPC header length is fixed at 80 bytes. Also, correct fragment order of fragmented connectionless DCE/RPC cannot be handled by a connectionless transport and, instead, must be ensured by connectionless DCE/RPC header values; in contrast, the transport protocol ensures correct fragment order for connection-oriented DCE/RPC. The DCE/RPC preprocessor uses these and other protocol-specific characteristics to monitor both protocols for anomalies and other evasion techniques, and to decode and defragment traffic before passing it to the rules engine.

The following diagram illustrates the point at which the DCE/RPC preprocessor begins processing DCE/RPC traffic for the different transports.



Note the following in the figure:

- The well-known TCP or UDP port 135 identifies DCE/RPC traffic in the TCP and UDP transports.
- The figure does not include RPC over HTTP.

  For RPC over HTTP, connection-oriented DCE/RPC is transported directly over TCP as shown in the figure after an initial setup sequence over HTTP. See Understanding the RPC over HTTP Transport, page 22-7 for more information.

- The DCE/RPC preprocessor typically receives SMB traffic on the well-known TCP port 139 for the NetBIOS Session Service or the similarly implemented well-known Windows port 445.

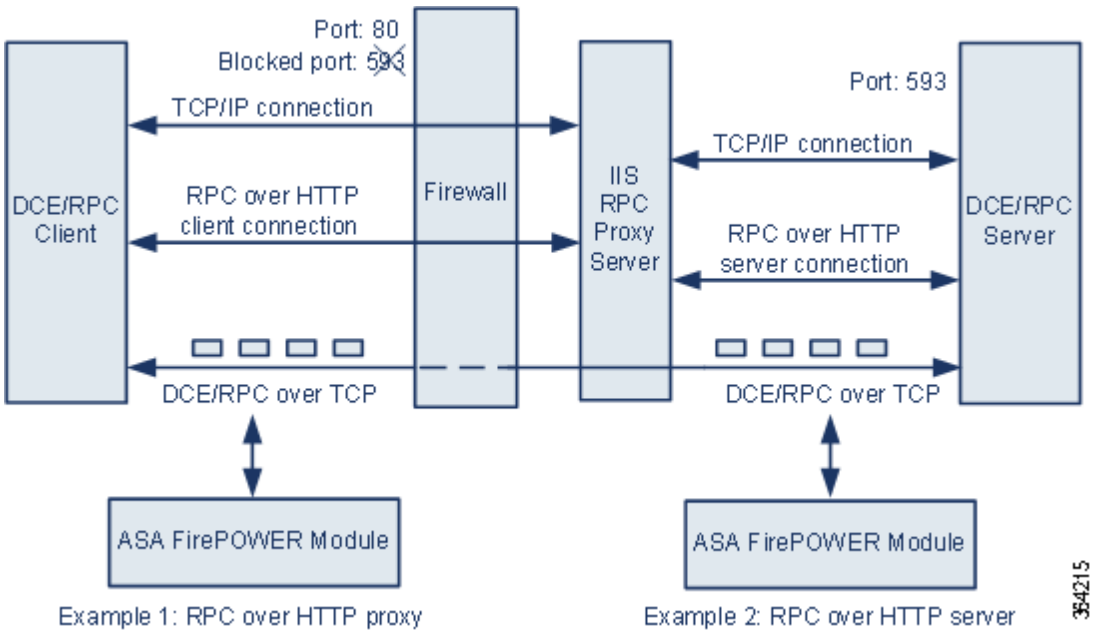
Because SMB has many functions other than transporting DCE/RPC, the preprocessor first tests whether the SMB traffic is carrying DCE/RPC traffic, stops processing if it is not, and continues processing if it is.

- IP encapsulates all DCE/RPC transports.
- TCP transports all connection-oriented DCE/RPC.
- UDP transports connectionless DCE/RPC.

Understanding the RPC over HTTP Transport

License: Protection

Microsoft RPC over HTTP allows you to tunnel DCE/RPC traffic through a firewall as shown in the following diagram. The DCE/RPC preprocessor detects version 1 of Microsoft RPC over HTTP.



The Microsoft IIS proxy server and the DCE/RPC server can be on the same host or on different hosts. Separate proxy and server options provide for both cases. Note the following in the figure:

- The DCE/RPC server monitors port 593 for DCE/RPC client traffic, but the firewall blocks port 593.

  Firewalls typically block port 593 by default.
- RPC over HTTP transports DCE/RPC over HTTP using well-known HTTP port 80, which firewalls are likely to permit.
- Example 1 shows that you would select the **RPC over HTTP proxy** option to monitor traffic between the DCE/RPC client and the Microsoft IIS RPC proxy server.
- Example 2 shows that you would select the **RPC over HTTP server** option when the Microsoft IIS RPC proxy server and the DCE/RPC server are located on different hosts and the device monitors traffic between the two servers.
- Traffic is comprised solely of connection-oriented DCE/RPC over TCP after RPC over HTTP completes the proxied setup between the DCE/RPC client and server.

Selecting DCE/RPC Target-Based Policy Options

License: Protection

Each target-based policy allows you to specify the various options below. Note that, except for the **Memory Cap Reached** and **Auto-Detect Policy on SMB Session** options, modifying these options could have a negative impact on performance or detection capability. You should not modify them unless you have a thorough understanding of the preprocessor and the interaction between the preprocessor and enabled DCE/RPC rules.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Networks

The host IP addresses where you want to apply the DCE/RPC target-based server policy.

You can specify a single IP address or address block, or a comma-separated list of either or both. You can specify up to 255 total profiles including the default policy. For information on specifying IPv4 and IPv6 address blocks in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or CIDR block/prefix length for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Note also that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks, and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

Policy

The Windows or Samba DCE/RPC implementation used by the targeted host or hosts on your monitored network segment. See Understanding Target-Based DCE/RPC Server Policies, page 22-4 for detailed information on these policies.

Note that you can enable the **Auto-Detect Policy on SMB Session** global option to automatically override the setting for this option on a per session basis when SMB is the DCE/RPC transport. See Auto-Detect Policy on SMB Session, page 22-3.

SMB Invalid Shares

A case-insensitive, alphanumeric text string that identifies one or more SMB shared resources; the preprocessor will detect when there is an attempt to connect to a shared resource that you specify. You can specify multiple shares in a comma-separated list and, optionally, you can enclose shares in quotes, which was required in previous software versions but is no longer required; for example:

```
"C$", D$, "admin", private
```

The preprocessor detects invalid shares in SMB traffic when you have enabled both SMB ports and SMB traffic detection.

Note that in most cases you should append a dollar sign to a drive named by Windows that you identify as an invalid share. For example, identify drive C as C$ or "C$".

You can enable rule 133:26 to generate events for this option. See Setting Rule States, page 27-19 for more information.

SMB Maximum AndX Chain

The maximum number between 0 and 255 of chained SMB AndX commands to permit. Typically, more than a few chained AndX commands represent anomalous behavior and could indicate an evasion attempt. Specify 1 to permit no chained commands or 0 to disable detecting the number of chained commands.

Note that the preprocessor first counts the number of chained commands and generates an event if accompanying SMB preprocessor rules are enabled and the number of chained commands equals or exceeds the configured value. It then continues processing.

Note Only someone who is expert in the SMB protocol should modify the default setting for this option.

You can enable rule 133:20 to generate events for this option. See Setting Rule States, page 27-19 for more information.

RPC proxy traffic only

When **RPC over HTTP Proxy Ports** is enabled, indicates whether detected client-side RPC over HTTP traffic is proxy traffic only or might include other web server traffic. For example, port 80 could carry both proxy and other web server traffic.

When this option is disabled, both proxy and other web server traffic are expected. Enable this option, for example, if the server is a dedicated proxy server. When enabled, the preprocessor tests traffic to determine if it carries DCE/RPC, ignores the traffic if it does not, and continues processing if it does. Note that enabling this option adds functionality only if the **RPC over HTTP Proxy Ports** check box is also enabled.

RPC over HTTP Proxy Ports

Enables detection of DCE/RPC traffic tunneled by RPC over HTTP over each specified port when your device is positioned between the DCE/RPC client and the Microsoft IIS RPC proxy server. See Understanding the RPC over HTTP Transport, page 22-7.

When enabled, you can add any ports where you see DCE/RPC traffic, although this is unlikely to be necessary because web servers typically use the default port for both DCE/RPC and other traffic. When enabled, you would not enable **RPC over HTTP Proxy Auto-Detect Ports**, but you would enable the **RPC Proxy Traffic Only** when detected client-side RPC over HTTP traffic is proxy traffic only and does not include other web server traffic.

RPC over HTTP Server Ports

Enables detection of DCE/RPC traffic tunneled by RPC over HTTP on each specified port when the Microsoft IIS RPC proxy server and the DCE/RPC server are located on different hosts and the device monitors traffic between the two servers. See Understanding the RPC over HTTP Transport, page 22-7.

Typically, when you enable this option you should also enable **RPC over HTTP Server Auto-Detect Ports** with a port range from 1025 to 65535 for that option even if you are not aware of any proxy web servers on your network. Note that the RPC over HTTP server port is sometimes reconfigured, in which case you should add the reconfigured server port to port list for this option.

TCP Ports

Enables detection of DCE/RPC traffic in TCP on each specified port.

Legitimate DCE/RPC traffic and exploits might use a wide variety of ports, and other ports above port 1024 are common. Typically, when this option is enabled you should also enable **TCP Auto-Detect Ports** with a port range from 1025 to 65535 for that option.

UDP Ports

Enables detection of DCE/RPC traffic in UDP on each specified port.

Legitimate DCE/RPC traffic and exploits might use a wide variety of ports, and other ports above port 1024 are common. Typically, when this option is enabled you should also enable **UDP Auto-Detect Ports** with a port range from 1025 to 65535 for that option.

SMB Ports

Enables detection of DCE/RPC traffic in SMB on each specified port.

You could encounter SMB traffic using the default detection ports. Other ports are rare. Typically, use the default settings.

RPC over HTTP Proxy Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic tunneled by RPC over HTTP on the specified ports when your device is positioned between the DCE/RPC client and the Microsoft IIS RPC proxy server. See Understanding the RPC over HTTP Transport, page 22-7.

When enabled, you would typically specify a port range from 1025 to 65535 to cover the entire range of ephemeral ports.

RPC over HTTP Server Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic tunneled by RPC over HTTP on the specified ports when the Microsoft IIS RPC proxy server and the DCE/RPC server are located on different hosts and the device monitors traffic between the two servers. See Understanding the RPC over HTTP Transport, page 22-7.

TCP Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic in TCP on the specified ports.

UDP Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic in UDP on each specified port.

SMB Auto-Detect Ports

Enables auto-detection of DCE/RPC traffic in SMB.

SMB File Inspection

Enables inspection of SMB traffic for file detection. You have the following options:

- Select **Off** to disable file inspection.
- Select **Only** to inspect file data without inspecting the DCE/RPC traffic in SMB. Selecting this option can improve performance over inspecting both files and DCE/RPC traffic.
- Select **On** to inspect both files and the DCE/RPC traffic in SMB. Selecting this option can impact performance.

Inspection of SMB traffic for the following is not supported:

- files transferred in SMB 2.x and SMB 3.x

- files transferred in an established TCP or SMB session before this option is enabled and the policy applied
- files transferred concurrently in a single TCP or SMB session
- files transferred across multiple TCP or SMB sessions
- files transferred with non-contiguous data, such as when message signing is negotiated
- files transferred with different data at the same offset, overlapping the data
- files opened on a remote client for editing that the client saves to the file server

SMB File Inspection Depth

If **SMB File Inspection** is set to **Only** or **On**, the number of bytes inspected when a file is detected in SMB traffic. Specify one of the following:

- an integer from `1` to `2147483647` (about 2GB)
- `0` to inspect the entire file
- `-1` to disable file inspection

Enter a value in this field equal to or smaller than the one defined in your access control policy. If you set a value for this option larger than the one defined for **Limit the number of bytes inspected when doing file type detection**, the system uses the access control policy setting as the functional maximum.See Tuning File and Malware Inspection Performance and Storage, page 11-16 for more information.

If **SMB File Inspection** is set to **Off**, this field is disabled.

Configuring the DCE/RPC Preprocessor

License: Protection

You can configure DCE/RPC preprocessor global options and one or more target-based server policies.

The preprocessor does not generate events unless you enable rules with generator ID (GID) 133. See Selecting Global DCE/RPC Options, page 22-3 and Selecting DCE/RPC Target-Based Policy Options, page 22-8 for rules associated with specific detection options; see also Setting Rule States, page 27-19.

In addition, most DCE/RPC preprocessor rules generate events against anomalies and evasion techniques detected in SMB, connection-oriented DCE/RPC, or connectionless DCE/RPC traffic. The following table identifies the rules that you can enable for each type of traffic.

Table 22-1 Traffic-Associated DCE/RPC Rules

Traffic	Preprocessor Rule GID:SID
SMB	133:2 through 133:26, and 133:48 through 133:57
Connection-Oriented DCE/RPC	133:27 through 133:39
Detect Connectionless DCE/RPC	133:40 through 133:43

To configure the DCE/RPC preprocessor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **DCE/RPC Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The DCE/RPC Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 You can modify any of the options described in Selecting Global DCE/RPC Options, page 22-3.

Step 10 You have two options:

- Add a new target-based policy. Click the add icon ( ) next to **Servers** on the left side of the page. The Add Target pop-up window appears. Specify a one or more IP addresses in the **Server Address** field and click **OK**.

  You can specify a single IP address or address block, or a comma-separated list of either or both. For information on using IPv4 and IPv6 address blocks in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

  You can configure up to 255 policies, including the default policy.

  Note that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks, and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

  A new entry appears in the list of servers on the left side of the page, highlighted to indicate that it is selected, and the Configuration section updates to reflect the current configuration for the profile you added.

- Modify the settings for an existing target-based policy. Click the configured address for a policy you have added under **Servers** on the left side of the page, or click **default**.

  Your selection is highlighted and the Configuration section updates to display the current configuration for the policy you selected. To delete an existing policy, click the delete icon ( ) next to the policy you want to remove.

Step 11 You can modify any of the following target-based policy options:

- To specify the host or hosts where you want to apply the DCE/RPC target-based server policy, enter a single IP address or address block, or a comma-separated list of either or both in the **Networks** field.

  You can specify up to 255 total profiles including the default policy. Note that you cannot modify the setting for **Networks** in the default policy. The default policy applies to all servers on your network that are not identified in another policy.

- To specify the type of policy you want to apply to the specified host or hosts on your network segment, select one of the Windows or Samba policy types from the **Policy** drop-down list.

  Note that you can enable the **Auto-Detect Policy on SMB Session** global option to automatically override the setting for this option on a per session basis when SMB is the DCE/RPC transport. See Auto-Detect Policy on SMB Session, page 22-3.

- To set the preprocessor to detect when there is an attempt to connect to specified shared SMB resources, enter a single or comma-separated list of the case-insensitive strings that identify the shared resources in the **SMB Invalid Shares** field. Optionally, enclose individual strings in quotes, which was required in previous software versions but is no longer required.

  For example, to detect shared resources named C$, D$, admin, and private, you could enter:

  ```
  "C$", D$, "admin", private
  ```

  Note that to detect SMB invalid shares, you must also enable **SMB Ports** or **SMB Auto-Detect Ports,** and enable the global **SMB Traffics** option.

  Note also that in most cases you should append a dollar sign to a drive named by Windows that you identify as an invalid share. For example, you would enter `C$` or `"C$"` to identify drive C.

- To inspect files detected in DCE/RPC traffic in SMB without analyzing the DCE/RPC traffic, from the **SMB File Inspection** drop-down list, select **Only**. To inspect files detected in DCE/RPC traffic in SMB as well as the DCE/RPC traffic, from the **SMB File Inspection** drop-down list, select **On**. Enter a number of bytes to inspect in a detected file in the **SMB File Inspection Depth** field. Enter `0` to inspect detected files in their entirety.

- To specify a maximum number of chained SMB AndX commands to permit, enter 0 to 255 in the **SMB Maximum AndX Chains** field. Specify `1` to permit no chained commands. Specify `0` or leave this option blank to disable this feature.

  Note Only someone who is expert in the SMB protocol should modify the setting for the **SMB Maximum AndX Chains** option.

- To enable the processing of DCE/RPC traffic over ports known to carry DCE/RPC traffic for a Windows policy transport, select or clear the check box next to a detection transport and, optionally, add or delete ports for the transport.

  Select one or any combination of **RPC over HTTP Proxy Ports**, **RPC over HTTP Server Ports**, **TCP Ports**, and **UDP Ports** for a Windows policy. Select **RPC Proxy Traffic Only** when **RPC over HTTP proxy** is enabled and detected client-side RPC over HTTP traffic is proxy traffic only; that is, when it does not include other web server traffic.

  Select **SMB Ports** for a Samba policy.

  In most cases, use the default settings. See Understanding DCE/RPC Transports, page 22-5, Understanding the RPC over HTTP Transport, page 22-7, and Selecting DCE/RPC Target-Based Policy Options, page 22-8 for more information.

  You can type a single port, a range of port numbers separated by a dash (-), or a comma-separated list of port numbers and ranges.

- To test whether specified ports carry DCE/RPC traffic and continue processing when they do, select or clear the check box next to an auto-detection transport and, optionally, add or delete ports for the transport.

  Select one or any combination of **RPC over HTTP Server Auto-Detect Ports**, **TCP Auto-Detect Ports**, and **UDP Auto-Detect Ports** for a Windows policy.

  Note that you would rarely, if ever, select **RPC over HTTP Proxy Auto-Detect Ports** or **SMB Auto-Detect Ports**.

  Typically, specify a port range from 1025 to 65535 for auto-detection ports that you enable to cover the entire range of ephemeral ports. See Understanding DCE/RPC Transports, page 22-5, Understanding the RPC over HTTP Transport, page 22-7, and Selecting DCE/RPC Target-Based Policy Options, page 22-8 for more information.

See Selecting DCE/RPC Target-Based Policy Options, page 22-8 for more information.

Step 12 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Detecting Exploits in DNS Name Server Responses

License: Protection

The DNS preprocessor inspects DNS name server responses for the following specific exploits:

- Overflow attempts on RData text fields
- Obsolete DNS resource record types
- Experimental DNS resource record types

See the following sections for more information:

- Understanding DNS Preprocessor Resource Record Inspection, page 22-14
- Detecting Overflow Attempts in RData Text Fields, page 22-15
- Detecting Obsolete DNS Resource Record Types, page 22-16
- Detecting Experimental DNS Resource Record Types, page 22-16
- Configuring the DNS Preprocessor, page 22-17

Understanding DNS Preprocessor Resource Record Inspection

License: Protection

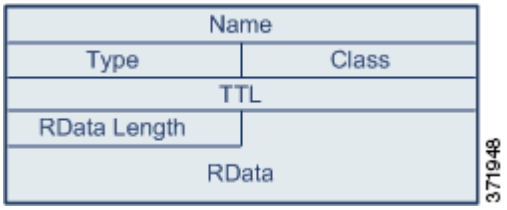
The most common type of DNS name server response provides one or more IP addresses that correspond to domain names in the query that prompted the response. Other types of server responses provide, for example, the destination for an email message or the location of a name server that can provide information not available from the server originally queried.

A DNS response is comprised of a message header, a Question section that contains one or more requests, and three sections that respond to requests in the Question section (Answer, Authority, and Additional Information). Responses in these three sections reflect the information in *resource records* (RR) maintained on the name server. The following table describes these three sections.

Table 22-2 DNS Name Server RR Responses

This section...	Includes...	For example...
Answer	Optionally, one or more resource records that provide a specific answer to a query	The IP address corresponding to a domain name
Authority	Optionally, one or more resource records that point to an authoritative name server	The name of an authoritative name server for the response
Additional Information	Optionally, one or more resource records that provided additional information related to the Answer sections	The IP address of another server to query

There are many types of resource records, all adhering to the following structure:



Theoretically, any type of resource record can be used in the Answer, Authority, or Additional Information section of a name server response message. The DNS preprocessor inspects any resource record in each of the three response sections for the exploits it detects.

The Type and RData resource record fields are of particular importance to the DNS preprocessor. The Type field identifies the type of resource record. The RData (resource data) field provides the response content. The size and content of the RData field differs depending on the type of resource record.

DNS messages typically use the UDP transport protocol but also use TCP when the message type requires reliable delivery or the message size exceeds UDP capabilities. The DNS preprocessor inspects DNS server responses in both UDP and TCP traffic.

The DNS preprocessor does not inspect TCP sessions picked up in midstream, and ceases inspection if a session loses state because of dropped packets.

The typical port to configure for the DNS preprocessor is well-known port 53, which DNS name servers use for DNS messages in both UDP and TCP.

Detecting Overflow Attempts in RData Text Fields

License: Protection

When the resource record type is TXT (text), the RData field is a variable-length ASCII text field.

When selected, the DNS preprocessor **Detect Overflow attempts on RData Text fields** option detects a specific vulnerability identified by entry CVE-2006-3441 in MITRE's Current Vulnerabilities and Exposures database. This is a known vulnerability in Microsoft Windows 2000 Service Pack 4, Windows XP Service Pack 1 and Service Pack 2, and Windows Server 2003 Service Pack 1. An attacker can

exploit this vulnerability and take complete control of a host by sending or otherwise causing the host to receive a maliciously crafted name server response that causes a miscalculation in the length of an RData text field, resulting in a buffer overflow.

You should enable this feature when your network might include hosts running operating systems that have not been upgraded to correct this vulnerability.

You can enable rule 131:3 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detecting Obsolete DNS Resource Record Types

License: Protection

RFC 1035 identifies several resource record types as obsolete. Because these are obsolete record types, some systems do not account for them and may be open to exploits. You would not expect to encounter these record types in normal DNS responses unless you have purposely configured your network to include them.

You can configure the system to detect known obsolete resource record types. The following table lists and describes these record types.

Table 22-3 Obsolete DNS Resource Record Types

RR Type	Code	Description
3	MD	a mail destination
4	MF	a mail forwarder

You can enable rule 131:1 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detecting Experimental DNS Resource Record Types

License: Protection

RFC 1035 identifies several resource record types as experimental. Because these are experimental record types, some systems do not account for them and may be open to exploits. You would not expect to encounter these record types in normal DNS responses unless you have purposely configured your network to include them.

You can configure the system to detect known experimental resource record types. The following table lists and describes these record types.

Table 22-4 Experimental DNS Resource Record Types

RR Type	Code	Description
7	MB	a mailbox domain name
8	MG	a mail group member
9	MR	a mail rename domain name
10	NUL	a null resource record

You can enable rule 131:2 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Configuring the DNS Preprocessor

License: Protection

Use the following procedure to configure the DNS preprocessor. For more information on configuring the options on this page, see Detecting Overflow Attempts in RData Text Fields, page 22-15, Detecting Obsolete DNS Resource Record Types, page 22-16, and Detecting Experimental DNS Resource Record Types, page 22-16.

To configure the DNS preprocessor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **DNS Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The DNS Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 Optionally, you can modify any of the following in the Settings area:

- Specify the source port or ports the DNS preprocessor should monitor for DNS server responses in the **Ports** field. Separate multiple ports with commas.
- Select the **Detect Overflow Attempts on RData Text fields** check box to enable detection of buffer overflow attempts in RData text fields.

- Select the **Detect Obsolete DNS RR Types** check box to enable detection of obsolete resource record types.
- Select the **Detect Experimental DNS RR Types** check box to detect experimental resource record types.

Step 10 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Decoding FTP and Telnet Traffic

License: Protection

The FTP/Telnet decoder analyzes FTP and telnet data streams, normalizing FTP and telnet commands before processing by the rules engine.

You must enable FTP and telnet preprocessor rules, which have generator IDs (GIDs) of 125 and 126, if you want these rules to generate events. See Setting Rule States, page 27-19 for more information.

For more information, see the following topics:

- Understanding Global FTP and Telnet Options, page 22-18
- Configuring Global FTP/Telnet Options, page 22-19
- Understanding Telnet Options, page 22-20
- Configuring Telnet Options, page 22-21
- Understanding Server-Level FTP Options, page 22-22
- Configuring Server-Level FTP Options, page 22-25
- Understanding Client-Level FTP Options, page 22-28
- Configuring Client-Level FTP Options, page 22-29

Understanding Global FTP and Telnet Options

License: Protection

You can set global options to determine whether the FTP/Telnet decoder performs stateful or stateless inspection of packets, whether the decoder detects encrypted FTP or telnet sessions, and whether the decoder continues to check a data stream after it encounters encrypted data.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Stateful Inspection

When selected, causes the FTP/Telnet decoder to save state and provide session context for individual packets and only inspects reassembled sessions. When cleared, analyzes each individual packet without session context.

To check for FTP data transfers, this option must be selected.

Detect Encrypted Traffic

Detects encrypted telnet and FTP sessions.

You can enable rules 125:7 and 126:2 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Continue to Inspect Encrypted Data

Instructs the preprocessor to continue checking a data stream after it is encrypted, looking for eventual decrypted data.

Configuring Global FTP/Telnet Options

License: Protection

You need to configure global options for the FTP/Telnet decoder to control whether stateless or stateful inspection is performed, encrypted traffic is detected, and whether the decoder should continue to check for decrypted data in a data stream that it has identified as encrypted. For more information on global settings, see Understanding Global FTP and Telnet Options, page 22-18.

To configure global options:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

The Advanced Settings page appears.

Step 8 You have two choices, depending on whether **FTP and Telnet Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The FTP and Telnet Configuration page appears.

A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.



Tip For more information on configuring the other options on this page, see Configuring Telnet Options, page 22-21, Configuring Server-Level FTP Options, page 22-25, and Configuring Client-Level FTP Options, page 22-29.

Step 9 Optionally, you can modify any of the following in the Global Settings page area:

- Select **Stateful Inspection** to examine reassembled TCP streams containing FTP packets. Clear **Stateful Inspection** to inspect only unreassembled packets.
- Select **Detect Encrypted Traffic** to detect encrypted traffic. Clear **Detect Encrypted Traffic** to ignore encrypted traffic.
- If needed, select **Continue to Inspect Encrypted Data** to continue checking a stream after it becomes encrypted, in case it becomes decrypted again and can be processed.

Step 10 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Understanding Telnet Options

License: Protection

You can enable or disable normalization of telnet commands by the FTP/Telnet decoder, enable or disable a specific anomaly case, and set the threshold number of Are You There (AYT) attacks to permit.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Ports

Indicates the ports whose telnet traffic you want to normalize. In the interface, list multiple ports separated by commas.

Normalize

Normalizes telnet traffic to the specified ports.

Detect Anomalies

Enables detection of Telnet SB (subnegotiation begin) without the corresponding SE (subnegotiation end).

Telnet supports subnegotiation, which begins with SB (subnegotiation begin) and must end with an SE (subnegotiation end). However, certain implementations of Telnet servers will ignore the SB without a corresponding SE. This is anomalous behavior that could be an evasion case. Because FTP uses the Telnet protocol on the control connection, it is also susceptible to this behavior.

You can enable rule 126:3 to generate an event when this anomaly is detected in Telnet traffic, and rule 125:9 when it is detected on the FTP command channel. See Setting Rule States, page 27-19 for more information.

Are You There Attack Threshold Number

Detects when the number of consecutive AYT commands exceeds the specified threshold. Cisco recommends that you set the AYT threshold to a value no higher than 20.

You can enable rule 126:1 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Configuring Telnet Options

License: Protection

You can enable or disable normalization, enable or disable a specific anomaly case, and control the threshold number of Are You There (AYT) attacks to permit. For additional information on telnet options, see Understanding Telnet Options, page 22-20.

To configure telnet options:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **FTP and Telnet Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The FTP and Telnet Configuration page appears.

A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Tip For more information on configuring the other options on this page, see Configuring Global FTP/Telnet Options, page 22-19, Configuring Server-Level FTP Options, page 22-25, and Configuring Client-Level FTP Options, page 22-29.

Step 9 Optionally, you can modify any of the following in the Telnet Settings page area:

- Specify the port or ports where telnet traffic should be decoded in the **Ports** field. Telnet typically connects to TCP port 23. Separate multiple ports with commas.

Caution Because encrypted traffic (SSL) cannot be decoded, adding port 22 (SSH) could yield unexpected results.

- Select or clear the **Normalize** Telnet Protocol Options check box to enable or disable telnet normalization.
- Select or clear the **Detect Anomalies** Telnet Protocol Options check box to enable or disable anomaly detection.
- Specify an **Are You There Attack Threshold Number** of consecutive AYT commands to permit.

Tip Cisco recommends that you set the AYT threshold to a value no higher than the default value.

Step 10 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Understanding Server-Level FTP Options

License: Protection

You can set options for decoding on multiple FTP servers. Each server profile you create contains the server IP address and the ports on the server where traffic should be monitored. You can specify which FTP commands to validate and which to ignore for a particular server, and set maximum parameter lengths for commands. You can also set the specific command syntax the decoder should validate against for particular commands and set alternate maximum command parameter lengths.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Networks

Use this option to specify one or more IP addresses of FTP servers.

You can specify a single IP address or address block, or a comma-separated list comprised of either or both. You can configure up to 1024 characters, and you can specify up to 255 profiles including the default profile. For information on using IPv4 and IPv6 address blocks in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or address block for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Note also that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks, and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

Ports

Use this option to specify the ports on the FTP server where the device should monitor traffic. In the interface, list multiple ports separated by commas.

File Get Commands

Use this option to define the FTP commands used to transfer files from server to client. Do not change these values unless directed to do so by Support.

File Put Commands

Use this option to define the FTP commands used to transfer files from client to server. Do not change these values unless directed to do so by Support.

Additional FTP Commands

Use this line to specify the additional commands that the decoder should detect. Separate additional commands by spaces.

Default Max Parameter Length

Use this option to detect the maximum parameter length for commands where an alternate maximum parameter length has not been set.

You can enable rule 125:3 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Alternate Max Parameter Length

Use this option to specify commands where you want to detect a different maximum parameter length, and to specify the maximum parameter length for those commands. Click **Add** to add lines where you can specify a different maximum parameter length to detect for particular commands.

Check Commands for String Format Attacks

Use this option to check the specified commands for string format attacks.

You can enable rule 125:5 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Command Validity

Use this option to enter a valid format for a specific command. See Creating FTP Command Parameter Validation Statements, page 22-24 for information on creating FTP command parameter validation statements to validate the syntax of a parameter received as part of an FTP communication. Click **Add** to add a command validation line.

You can enable rules 125:2 and 125:4 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Ignore FTP Transfers

Use this option to improve performance on FTP data transfers by disabling all inspection other than state inspection on the data transfer channel.

Detect Telnet Escape Codes within FTP Commands

Use this option to detect when telnet commands are used over the FTP command channel.

You can enable rule 125:1 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Ignore Erase Commands during Normalization

When **Detect Telnet Escape Codes within FTP Commands** is selected, use this option to ignore telnet character and line erase commands when normalizing FTP traffic. The setting should match how the FTP server handles telnet erase commands. Note that newer FTP servers typically ignore telnet erase commands, while older servers typically process them.

Troubleshooting Options: Log FTP Command Validation Configuration

Support might ask you during a troubleshooting call to configure your system to print the configuration information for each FTP command listed for the server.

Caution Changing the setting for this troubleshooting option affects performance and should be done only with Support guidance.

Creating FTP Command Parameter Validation Statements

License: Protection

When setting up a validation statement for an FTP command, you can specify a group of alternative parameters by separating the parameters with spaces. You can also create a binary OR relationship between two parameters by separating them with a pipe character (`|`) in the validation statement. Surrounding parameters by square brackets (`[]`) indicates that those parameters are optional. Surrounding parameters with curly brackets (`{}`) indicates that those parameters are required.

You can create FTP command parameter validation statements to validate the syntax of a parameter received as part of an FTP communication. See Understanding Server-Level FTP Options, page 22-22 for more information.

Any of the parameters listed in the following table can be used in FTP command parameter validation statements.

Table 22-5 FTP Command Parameters

If you use...	The following validation occurs...
`int`	The represented parameter must be an integer.
`number`	The represented parameter must be an integer between 1 and 255.
`char _chars`	The represented parameter must be a single character and a member of the characters specified in the `_chars` argument. For example, defining the command validity for `MODE` with the validation statement `char SBC` checks that the parameter for the `MODE` command comprises the character `S` (representing Stream mode), the character `B` (representing Block mode), or the character `C` (representing Compressed mode).

Table 22-5 FTP Command Parameters (continued)

If you use...	The following validation occurs...
`date _datefmt`	If `_datefmt` contains `#`, the represented parameter must be a number. If `_datefmt` contains `C`, the represented parameter must be a character. If `_datefmt` contains literal strings, the represented parameter must match the literal string.
`string`	The represented parameter must be a string.
`host_port`	The represented parameter must be a valid host port specifier as defined by RFC 959, the File Transfer Protocol specification by the Network Working Group.

You can combine the syntax in the table above as needed to create parameter validation statements that correctly validate each FTP command where you need to validate traffic.

Note When you include a complex expression in a TYPE command, surround it by spaces. Also, surround each operand within the expression by spaces. For example, type `char A | B` , not `char A|B`.

Configuring Server-Level FTP Options

License: Protection

You can configure several options at the server level. For each FTP server you add, you can specify the ports to be monitored, the commands to validate, the default maximum parameter length for commands, alternate parameter lengths for specific commands, and validation syntax for particular commands. You can also choose whether to check for string format attacks and telnet commands on the FTP channel and whether to print configuration information with each command. For additional information on server-level FTP options, see Understanding Server-Level FTP Options, page 22-22.

To configure server-level FTP options:

Step 11 Click the edit icon ( ) next to the policy you want to edit.

Step 1 If you have unsaved changes Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 6 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 7 You have two choices, depending on whether **FTP and Telnet Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The FTP and Telnet Configuration page appears.

A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Tip For more information on configuring the other options on this page, see Configuring Global FTP/Telnet Options, page 22-19, Configuring Telnet Options, page 22-21, and Configuring Client-Level FTP Options, page 22-29.

Step 8 You have two options:

- Add a new server profile. Click the add icon ( ) next to **FTP Server** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses for the client in the **Server Address** field and click **OK**.

  You can specify a single IP address or address block, or a comma-separated list of either or both. You can specify up to 1024 characters, and you can configure up to 255 policies, including the default policy. For information on using IPv4 and IPv6 address blocks in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

  Note that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks, and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

  A new entry appears in the list of FTP servers on the left side of the page, highlighted to indicate that it is selected, and the Configuration section updates to reflect the current configuration for the profile you added.

- Modify the settings for an existing server profile. Click the configured address for a profile you have added under **FTP Server** on the left side of the page, or click **default**.

  Your selection is highlighted and the Configuration section updates to display the current configuration for the profile you selected. To delete an existing profile, click the delete icon ( ) next to the profile you want to remove.

Step 9 Optionally, you can modify any of the following in the Configuration page area:

- Modify the address or addresses listed in the **Networks** field and click any other area of the page.

  The highlighted address updates on the left side of the page.

  Note that you cannot modify the setting for **Network** in the default profile. The default profile applies to all servers on your network that are not identified in another profile.

- Specify any **Ports** that should be monitored for FTP traffic. Port 21 is the well-known port for FTP traffic.
- Update the FTP commands used to transfer files from server to client in the **File Get Commands** field.

- Update the FTP commands used to transfer files from client to server in the **File Put Commands** field.

Note Do not change the values in the **File Get Commands** and **File Put Commands** field unless directed to do so by Support.

- To detect additional FTP commands outside of those checked by default by the FTP/Telnet preprocessor, type the commands, separated by spaces in the **Additional FTP Commands** field.

  You can add as many additional FTP commands as needed.

Note Additional commands you may want to add include `XPWD`, `XCWD`, `XCUP`, `XMKD`, and `XRMD`. For more information on these commands, see RFC 775, the Directory oriented FTP commands specification by the Network Working Group.

- Specify the default maximum number of bytes for a command parameter in the **Default Max Parameter Length** field.
- To detect a different maximum parameter length for particular commands, click **Add** next to **Alternate Max Parameter Length**. In the first text box of the row that appears, specify the maximum parameter length. In the second text box, specify the commands, separated by spaces, where this alternate maximum parameter length should apply.

  You can add as many alternative maximum parameter lengths as needed.
- To check for string format attacks on particular commands, specify the commands, separated by spaces, in the **Check Commands for String Format Attacks** text box.
- To specify the valid format for a command, click **Add** next to **Command Validity**. Specify the command you want to validate, then type a validation statement for the command parameter. For more information on the validation statement syntax, see Understanding Server-Level FTP Options, page 22-22.
- To improve performance on FTP data transfers by disabling all inspection other than state inspection on the data transfer channel, enable **Ignore FTP Transfers**.

Note To inspect data transfers, the global FTP/Telnet **Stateful Inspection** option must be selected. For more information on setting global options, see Understanding Global FTP and Telnet Options, page 22-18.

- To detect when telnet commands are used over the FTP command channel, select **Detect Telnet Escape Codes within FTP Commands**.
- To ignore telnet character and line erase commands when normalizing FTP traffic, enable **Ignore Erase Commands during Normalization**.

Step 10 Optionally, modify the related troubleshooting option only if asked to do so by Support; click the + sign next to **Troubleshooting Options** to expand the troubleshooting options section.

Step 11 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Understanding Client-Level FTP Options

License: Protection

Use these options to configure custom FTP client profiles. If an option description does not include a preprocessor rule, the option is not associated with a preprocessor rule.

Networks

Use this option to specify one or more IP addresses of FTP clients.

You can specify a single IP address or address block, or a comma-separated list comprised of either or both. You can specify up to 1024 characters, and you can specify up to 255 profiles including the default profile. For information on using IPv4 and IPv6 address blocks in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or address block for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Note also that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks, and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

Max Response Length

Use this option to specify the maximum allowed response length to an FTP command accepted by the FTP client. It can be used as a basic buffer overflow detection.

You can enable rule 125:6 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect FTP Bounce Attempts

Use this option to detect FTP bounce attacks.

You can enable rule 125:8 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Allow FTP Bounce to

Use this option to configure a list of additional hosts and ports on those hosts on which FTP PORT commands should not be treated as FTP bounce attacks.

Detect Telnet Escape Codes within FTP Commands

Use this option to detect when telnet commands are used over the FTP command channel.

You can enable rule 125:1 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Ignore Erase Commands During Normalization

When **Detect Telnet Escape Codes within FTP Commands** is selected, use this option to ignore telnet character and line erase commands when normalizing FTP traffic.The setting should match how the FTP client handles telnet erase commands. Note that newer FTP clients typically ignore telnet erase commands, while older clients typically process them.

Configuring Client-Level FTP Options

License: Protection

You can configure client profiles for FTP clients to monitor FTP traffic from clients. For additional information on the options you can set for monitoring clients, see Understanding Client-Level FTP Options, page 22-28. For more information on telnet options, see Understanding Telnet Options, page 22-20. For more information on additional FTP options, see Understanding Server-Level FTP Options, page 22-22 and Understanding Global FTP and Telnet Options, page 22-18.

To configure client-level FTP options:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **FTP and Telnet Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The FTP and Telnet Configuration page appears.

Step 9 You have two options:

- Add a new client profile. Click the add icon ( ) next to **FTP Client** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses for the client in the **Client Address** field and click **OK**.

  You can specify a single IP address or address block, or a comma-separated list of either or both. You can specify up to 1024 characters, and you can configure up to 255 policies, including the default policy. For information on using IPv4 and IPv6 address blocks in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

Note that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks, and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

A new entry appears in the list of FTP clients on the left side of the page, highlighted to indicate that it is selected, and the Configuration section updates to reflect the current configuration for the profile you added.

- Modify the settings for an existing client profile. Click the configured address for a profile you have added under **FTP Client on** the left side of the page, or click **default**.

  Your selection is highlighted and the Configuration section updates to display the current configuration for the profile you selected. To delete an existing profile, click the delete icon () next to the profile you want to remove.

Step 10 Optionally, you can modify any of the following in the Configuration page area:

- Optionally, modify the address or addresses listed in the **Networks** field and click any other area of the page.

  The highlighted address updates on the left side of the page.

  Note that you cannot modify the setting for **Network** in the default profile. The default profile applies to all client hosts on your network that are not identified in another profile.

- Specify, in bytes, the maximum length of responses from the FTP client in the **Max Response Length** field.

- To detect FTP bounce attacks, select **Detect FTP Bounce attempts**.

  The FTP/Telnet decoder detects when an FTP PORT command is issued and the specified host does not match the specified host of the client.

- To configure a list of additional hosts and ports where FTP PORT commands should not be treated as FTP bounce attacks, specify each host (or network in CIDR format) followed by a colon (:) and the port or port range in the **Allow FTP Bounce to** field. To enter a range of ports for a host, separate the beginning port in the range and the final port in the range with a dash (-). You can enter multiple hosts by separating the entries for the hosts with a comma.

  For example, to permit FTP PORT commands directed to the host 192.168.1.1 at port 21 and commands directed to the host 192.168.1.2 at any of the ports from 22 to 1024, type:

  ```
  192.168.1.1:21, 192.168.1.2:22-1024
  ```

  For information on using CIDR notation and prefix lengths in the ASA FirePOWER module, see IP Address Conventions, page 1-4.



Note To specify multiple individual ports for a host, you must repeat the host IP address for each port definition. For example, to specify the ports 22 and 25 on 192.168.1.1, type `192.168.1.1:22, 192.168.1.1:25`.

- To detect when telnet commands are used over the FTP command channel, select **Detect Telnet Escape Codes within FTP Commands**.

- To ignore telnet character and line erase commands when normalizing FTP traffic, select **Ignore Erase Commands During Normalization**.

Step 11 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Decoding HTTP Traffic

License: Protection

The HTTP Inspect preprocessor is responsible for:

- decoding and normalizing HTTP requests sent to and HTTP responses received from web servers on your network
- separating messages sent to web servers into URI, non-cookie header, cookie header, method, and message body components to improve performance of HTTP-related intrusion rules
- separating messages received from web servers into status code, status message, non-set-cookie header, cookie header, and response body components to improve performance of HTTP-related intrusion rules
- detecting possible URI-encoding attacks
- making the normalized data available for additional rule processing

HTTP traffic can be encoded in a variety of formats, making it difficult for rules to appropriately inspect. HTTP Inspect decodes 14 types of encoding, ensuring that your HTTP traffic gets the best inspection possible.

You can configure HTTP Inspect options globally, on a single server, or for a list of servers.

Note the following when using the HTTP Inspect preprocessor:

- The preprocessor engine performs HTTP normalization *statelessly*. That is, it normalizes HTTP strings on a packet-by-packet basis, and can only process HTTP strings that have been reassembled by the TCP stream preprocessor.
- You must enable HTTP preprocessor rules, which have a generator ID (GID) of 119, if you want these rules to generate events. See Setting Rule States, page 27-19 for more information.

See the following sections for more information:

- Selecting Global HTTP Normalization Options, page 22-31
- Configuring Global HTTP Configuration Options, page 22-32
- Selecting Server-Level HTTP Normalization Options, page 22-33
- Selecting Server-Level HTTP Normalization Encoding Options, page 22-41
- Configuring HTTP Server Options, page 22-43
- Enabling Additional HTTP Inspect Preprocessor Rules, page 22-45

Selecting Global HTTP Normalization Options

License: Protection

The global HTTP options provided for the HTTP Inspect preprocessor control how the preprocessor functions. Use these options to enable or disable HTTP normalization when ports not specified as web server ports receive HTTP traffic.

Note the following:

- If you enable **Unlimited Decompression**, the **Maximum Compressed Data Depth** and **Maximum Decompressed Data Depth** options are automatically set to 65535 when you commit your changes. See Selecting Server-Level HTTP Normalization Options, page 22-33 for more information.
- If the values for the **Maximum Compressed Data Depth** and **Maximum Decompressed Data Depth** options are different in an intrusion policy associated with the default action of an access control policy and intrusion policies associated with access control rules, the highest value is used.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Detect Anomalous HTTP Servers

Detects HTTP traffic sent to or received by ports not specified as web server ports.



Note If you turn this option on, be make to list all ports that do receive HTTP traffic in a server profile on the HTTP Configuration page. If you do not, and you enable this option and the accompanying preprocessor rule, normal traffic to and from the server will generate events. The default server profile contains all ports normally used for HTTP traffic, but if you modified that profile, you may need to add those ports to another profile to prevent events from being generated.

You can enable rule 120:1 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect HTTP Proxy Servers

Detects HTTP traffic using proxy servers not defined by the **Allow HTTP Proxy Use** option.

You can enable rule 119:17 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Maximum Compressed Data Depth

Sets the maximum size of compressed data to decompress when **Inspect Compressed Data** (and, optionally, **Decompress SWF File (LZMA)**, **Decompress SWF File (Deflate)**, or **Decompress PDF File (Deflate)**) is enabled. You can specify from 1 to 65535 bytes.

Maximum Decompressed Data Depth

Sets the maximum size of the normalized decompressed data when **Inspect Compressed Data** (and, optionally, **Decompress SWF File (LZMA)**, **Decompress SWF File (Deflate)**, or **Decompress PDF File (Deflate)**) is enabled. You can specify from 1 to 65535 bytes.

Configuring Global HTTP Configuration Options

License: Protection

You can configure detection of HTTP traffic to non-standard ports and on HTTP traffic using proxy servers. For more information on global HTTP configuration options, see Selecting Global HTTP Normalization Options, page 22-31.

To configure global HTTP configuration options:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **HTTP Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The HTTP Configuration page appears.

Step 9 You can modify any of the global options described in Selecting Global HTTP Normalization Options, page 22-31.

Step 10 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Selecting Server-Level HTTP Normalization Options

License: Protection

You can set server-level options for each server you monitor, globally for all servers, or for a list of servers. Additionally, you can use a predefined server profile to set these options, or you can set them individually to meet the needs of your environment. Use these options, or one of the default profiles that set these options, to specify the HTTP server ports whose traffic you want to normalize, the amount of server response payload you want to normalize, and the types of encoding you want to normalize.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Networks

Use this option to specify the IP address of one or more servers. You can specify a single IP address or address block, or a comma-separated list comprised of either or both.

In addition to a limit of up to 255 total profiles, including the default profile, you can include up to 496 characters, or approximately 26 entries, in an HTTP server list, and specify a total of 256 address entries for all server profiles. For information on using IPv4 CIDR notation and IPv6 prefix lengths in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or CIDR block/prefix length for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Note also that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks, and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

Ports

The ports whose HTTP traffic the preprocessor engine normalizes. Separate multiple port numbers with commas.

Oversize Dir Length

Detects URL directories longer than the specified value.

You can enable rule 119:15 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Client Flow Depth

Specifies the number of bytes for rules to inspect in raw HTTP packets, including header and payload data, in client-side HTTP traffic defined in **Ports**. Client flow depth does not apply when HTTP content rule options within a rule inspect specific parts of a request message. See HTTP Content Options, page 30-23 for more information.

You can specify a value from -1 to 1460. Cisco recommends that you set client flow depth to its maximum value. Specify any of the following:

– From 1 to 1460 inspects the specified number of bytes in the first packet. If the first packet contains fewer bytes than specified, inspect the entire packet. Note that the specified value applies to both segmented and reassembled packets.

  Note also that a value of 300 typically eliminates inspection of large HTTP Cookies that appear at the end of many client request headers.

– 0 inspects all client-side traffic, including multiple packets in a session and exceeding the 1460 byte limit if necessary. Note that this value is likely to affect performance.

– -1 ignores all client-side traffic.

Server Flow Depth

Specifies the number of bytes for rules to inspect in raw HTTP packets in server-side HTTP traffic specified by **Ports**. Inspection includes the raw header and payload when **Inspect HTTP Responses** disabled and only the raw response body when **Inspect HTTP Response** is enabled.

Server flow depth specifies the number of bytes of raw server response data in a session for rules to inspect in server-side HTTP traffic defined in **Ports**. You can use this option to balance performance and the level of inspection of HTTP server response data. Server flow depth does not apply when HTTP content options within a rule inspect specific parts of a response message. See HTTP Content Options, page 30-23 for more information.

Unlike client flow depth, server flow depth specifies the number of bytes per HTTP response, not per HTTP request packet, for rules to inspect.

You can specify a value from -1 to 65535. Cisco recommends that you set the server flow depth to its maximum value. You can specify any of the following:

- From 1 to 65535:

  When **Inspect HTTP Responses** is **enabled**, inspects only the raw HTTP response body, and not raw HTTP headers; also inspects decompressed data when **Inspect Compressed Data** is enabled.

  When **Inspect HTTP Responses** is **disabled**, inspects the raw packet header and payload.

  If the session includes fewer response bytes than specified, rules fully inspect all response packets in a given session, across multiple packets as needed. If the session includes more response bytes than specified, rules inspect only the specified number of bytes for that session, across multiple packets as needed.

  Note that a small flow depth value may cause false negatives from rules that target server-side traffic defined in **Ports**. Most of these rules target either the HTTP header or content that is likely to be in the first hundred or so bytes of non-header data. Headers are usually under 300 bytes long, but header size may vary.

  Note also that the specified value applies to both segmented and reassembled packets.

- 0 inspects the entire packet for all HTTP server-side traffic defined in **Ports**, including response data in a session that exceeds 65535 bytes.

  Note that this value is likely to affect performance.

- -1:

  When **Inspect HTTP Responses** is **enabled**, inspects only raw HTTP headers and not the raw HTTP response body.

  When **Inspect HTTP Responses** is **disabled**, ignores all server-side traffic defined in **Ports.**

Maximum Header Length

Detects a header field longer than the specified maximum number of bytes in an HTTP request; also in HTTP responses when **Inspect HTTP Responses** is enabled. The value of 0 disables this option. Specify a value from 1 to 65535 to enable it.

You can enable rule 119:19 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Maximum Number of Headers

Detects when the number of headers exceeds this setting in an HTTP request. Specify a value from 1 to 1024 to enable it.

You can enable rule 119:20 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Maximum Number of Spaces

Detects when the number of white spaces in a folded line equals or exceeds this setting in an HTTP request. A value of 0 disables this option. Specify a value from 1 to 65535 to enable it.

You can enable rule 119:26 to generate events for this option. See Setting Rule States, page 27-19 for more information.

HTTP Client Body Extraction Depth

Specifies the number of bytes to extract from the message body of an HTTP client request. You can use an intrusion rule to inspect the extracted data by selecting the `content` or `protected_content` keyword **HTTP Client Body** option. See HTTP Content Options, page 30-23 for more information.

Specify a value from -1 to 65495. Specify -1 to ignore the client body. Specify 0 to extract the entire client body. Note that identifying specific bytes to extract can improve system performance. Note also that you must specify a value from 0 to 65495 for the **HTTP Client Body** option to function in an intrusion rule.

Small Chunk Size

Specifies the maximum number of bytes at which a chunk is considered small. Specify a value of 1 to 255. A value of 0 disables detection of anomalous consecutive small segments. See the **Consecutive Small Chunks** option for more information.

Consecutive Small Chunks

Specifies how many consecutive small chunks represent an abnormally large number in client or server traffic that uses chunked transfer encoding. The **Small Chunk Size** option specifies the maximum size of a small chunk.

For example, set **Small Chunk Size** to 10 and **Consecutive Small Chunks** to 5 to detect 5 consecutive chunks of 10 bytes or less.

You can enable preprocessor rule 119:27 to trigger events on excessive small chunks in client traffic, and rule 120:7 in server traffic. When **Small Chunk Size** is enabled and this option is set to 0 or 1, enabling these rules would trigger an event on every chunk of the specified size or less. See Setting Rule States, page 27-19 for more information.

HTTP Methods

Specifies HTTP request methods in addition to GET and POST that you expect the system to encounter in traffic. Use a comma to separate multiple values.

Intrusion rules use the `content` or `protected_content` keyword with the **HTTP Method** argument to search for content in HTTP methods. See HTTP Content Options, page 30-23. You can enable rule 119:31 to generate events when a method other than GET, POST, or a method configured for this option is encountered in traffic.

No Alerts

Disables intrusion events when accompanying preprocessor rules are enabled.



Note This option does **not** disable HTTP standard text rules and shared object rules.

Normalize HTTP Headers

When **Inspect HTTP Responses** is enabled, enables normalization of non-cookie data in request and response headers. When **Inspect HTTP Responses** is **not** enabled, enables normalization of the entire HTTP header, including cookies, in request and response headers.

Inspect HTTP Cookies

Enables extraction of cookies from HTTP request headers. Also enables extraction of set-cookie data from response headers when **Inspect HTTP Responses** is enabled. Disabling this option when cookie extraction is not required can improve performance.

Note that the `Cookie:` and `Set-Cookie:` header names, leading spaces on the header line, and the `CRLF` that terminates the header line are inspected as part of the header and not as part of the cookie.

Normalize Cookies in HTTP headers

Enables normalization of cookies in HTTP request headers. When **Inspect HTTP Responses** is enabled, also enables normalization of set-cookie data in response headers. You must select **Inspect HTTP Cookies** before selecting this options.

Allow HTTP Proxy Use

Allows the monitored web server to be used as an HTTP proxy. This option is used only in the inspection of HTTP requests.

Inspect URI Only

Inspects only the URI portion of the normalized HTTP request packet.

Inspect HTTP Responses

Enables extended inspection of HTTP responses so, in addition to decoding and normalizing HTTP request messages, the preprocessor extracts response fields for inspection by the rules engine. Enabling this option causes the system to extract the response header, body, status code, and so on, and also extracts set-cookie data when **Inspect HTTP Cookies** is enabled. For more information, see HTTP Content Options, page 30-23, Generating Events on the HTTP Encoding Type and Location, page 30-94, and Pointing to a Specific Payload Type, page 30-98.

You can enable rules 120:2 and 120:3 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Normalize UTF Encodings to UTF-8

When **Inspect HTTP Responses** is enabled, detects UTF-16LE, UTF-16BE, UTF-32LE, and UTF32-BE encodings in HTTP responses and normalizes them to UTF-8.

You can enable rule 120:4 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Inspect Compressed Data

When **Inspect HTTP Responses** is enabled, enables decompression of gzip and deflate-compatible compressed data in the HTTP response body, and inspection of the normalized decompressed data. The system inspects chunked and non-chunked HTTP response data. The system inspects decompressed data packet by packet across multiple packets as needed; that is, the system does not combine the decompressed data from different packets for inspection. Decompression ends when **Maximum Compressed Data Depth**, **Maximum Decompressed Data Depth**, or the end of the compressed data is reached. Inspection of decompressed data ends when **Server Flow Depth** is reached unless you also select **Unlimited Decompression**. You can use the `file_data` rule keyword to inspect decompressed data; see Pointing to a Specific Payload Type, page 30-98 for more information.

Unlimited Decompression

When **Inspect Compressed Data** (and, optionally, **Decompress SWF File (LZMA)**, **Decompress SWF File (Deflate)**, or **Decompress PDF File (Deflate)**) is enabled, overrides **Maximum Decompressed Data Depth** across multiple packets; that is, this option enables unlimited decompression across multiple packets. Note that enabling this option does not affect **Maximum Compressed Data Depth** or **Maximum Decompressed Data Depth** within a single packet. Note also that enabling this option sets **Maximum Compressed Data Depth** and **Maximum Decompressed Data Depth** to 65535 when you commit your changes. See Selecting Global HTTP Normalization Options, page 22-31.

Normalize Javascript

When **Inspect HTTP Responses** is enabled, enables detection and normalization of Javascript within the HTTP response body. The preprocessor normalizes obfuscated Javascript data such as the unescape and decodeURI functions and the String.fromCharCode method. The preprocessor normalizes the following encodings within the unescape, decodeURI, and decodeURIComponent functions:

- %XX
- %uXXXX
- 0xXX
- \xXX
- \uXXXX

The preprocessor detects consecutive white spaces and normalizes them into a single space. When this option is enabled, a configuration field allows you to specify the maximum number of consecutive white spaces to permit in obfuscated Javascript data. You can enter a value from 1 to 65535. The value 0 disables event generation, regardless of whether the preprocessor rule (120:10) associated with this field is enabled.

The preprocessor also normalizes the Javascript plus (+) operator and concatenates strings using the operator.

You can use the `file_data` keyword to point intrusion rules to the normalized Javascript data. See Pointing to a Specific Payload Type, page 30-98 for more information.

You can enable rules 120:9, 120:10, and 120:11 to generate events for this option, as follows:

Table 22-6 Normalize Javascript Option Rules

This rule...	Triggers an event when...
120:9	the obfuscation level within the preprocessor is greater than or equal to 2.
120:10	the number of consecutive white spaces in the Javascript obfuscated data is greater than or equal to the value configured for the maximum number of consecutive white spaces allowed.
120:11	escaped or encoded data includes more than one type of encoding.

See Setting Rule States, page 27-19 for more information.

Decompress SWF File (LZMA) and Decompress SWF File (Deflate)

When **HTTP Inspect Responses** is enabled, these options decompress the compressed portions of files located within the HTTP response body of HTTP requests.

Note You can **only** decompress the compressed portions of files found in HTTP GET responses.

- **Decompress SWF File (LZMA)** decompresses the LZMA-compatible compressed portions of Adobe ShockWave Flash (`.swf`) files
- **Decompress SWF File (Deflate)** decompresses the deflate-compatible compressed portions of Adobe ShockWave Flash (`.swf`) files

Decompression ends when **Maximum Compressed Data Depth**, **Maximum Decompressed Data Depth**, or the end of the compressed data is reached. Inspection of decompressed data ends when **Server Flow Depth** is reached unless you also select **Unlimited Decompression**. You can use the `file_data` rule keyword to inspect decompressed data; see Pointing to a Specific Payload Type, page 30-98 for more information.

You can enable rules 120:12 and 120:13 to generate events for this option, as follows:

Table 22-7 Decompress SWF File Option Rules

This rule...	Triggers an event when...
120:12	deflate file decompression fails.
120:13	LZMA file decompression fails.

Decompress PDF File (Deflate)

When **HTTP Inspect Responses** is enabled, **Decompress PDF File (Deflate)** decompresses the deflate-compatible compressed portions of Portable Document Format (`.pdf`) files located within the HTTP response body of HTTP requests. The system can only decompress PDF files with the `/FlateDecode` stream filter. Other stream filters (including `/FlateDecode /FlateDecode`) are unsupported.

Note You can **only** decompress the compressed portions of files found in HTTP GET responses.

Decompression ends when **Maximum Compressed Data Depth**, **Maximum Decompressed Data Depth**, or the end of the compressed data is reached. Inspection of decompressed data ends when **Server Flow Depth** is reached unless you also select **Unlimited Decompression**. You can use the `file_data` rule keyword to inspect decompressed data; see Pointing to a Specific Payload Type, page 30-98 for more information.

You can enable rules 120:14, 120:15, 120:16, and 120:17 to generate events for this option, as follows:

Table 22-8 Decompress PDF File (Deflate) Option Rules

This rule...	Triggers an event when...
120:14	file decompression fails.
120:15	file decompression fails due to an unsupported compression type.
120:16	file decompression fails due to an unsupported PDF stream filter.
120:17	file parsing fails.

Extract Original Client IP Address

Enables the examination of original client IP addresses during intrusion inspection. The system extracts the original client IP address from the X-Forwarded-For (XFF), True-Client-IP, or custom HTTP headers you define in the XFF Header Priority option. You can view the extracted original client IP address in the intrusion events table.

You can enable rules 119:23, 119:29 and 119:30 to generate intrusion events for this option.

XFF Header Priority

If **Extract Original Client IP Address** is enabled, specifies the order in which the system processes original client IP headers when multiple headers are present in an HTTP request. By default, the system examines X-Forwarded-For (XFF) headers, then True-Client-IP headers. Use the up and down arrow icons beside each header type to adjust its priority.

This option also allows you to specify original client IP headers other than XFF or True-Client-IP for extraction and evaluation. Click **Add** to add custom header names to the priority list. The system only supports custom headers that use the same syntax as an XFF or True-Client-IP header.

Keep in mind the following when configuring this option:

- The system uses this priority order when evaluating original client IP address headers for both access control and intrusion inspection.
- If multiple original client IP headers are present, the system processes only the header with the highest priority.
- The XFF header contains a list of IP addresses, which represent the proxy servers through which the request has passed. To prevent spoofing, the system uses the last IP address in the list (that is, the address appended by the trusted proxy) as the original client IP address.

Log URI

Enables extraction of the raw URI, if present, from HTTP request packets and associates the URI with all intrusion events generated for the session.

When this option is enabled, you can display the first fifty characters of the extracted URI in the HTTP URI column of the intrusion events table view. You can display the complete URI, up to 2048 bytes, in the packet view. See Viewing Events, page 37-1 for more information.

Log Hostname

Enables extraction of the host name, if present, from the HTTP request Host header and associates the host name with all intrusion events generated for the session. When multiple Host headers are present, extracts the host name from the first header.

When this option is enabled, you can display the first fifty characters of the extracted host name in the HTTP Hostname column of the intrusion events table view. You can display the complete host name, up to 256 bytes, in the packet view. See Viewing Events, page 37-1 for more information.

You can enable rule 119:25 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Note that when the preprocessor and rule 119:24 are enabled, the preprocessor generates an intrusion event if it detects multiple Host headers in an HTTP request, regardless of the setting for this option. See Enabling Additional HTTP Inspect Preprocessor Rules, page 22-45 for more information.

Profile

Specifies the types of encoding that are normalized for HTTP traffic. The system provides a default profile appropriate for most servers, default profiles for Apache servers and IIS servers, and custom default settings that you can tailor to meet the needs of your monitored traffic. See Selecting Server-Level HTTP Normalization Encoding Options, page 22-41 for more information.

Selecting Server-Level HTTP Normalization Encoding Options

License: Protection

You can select server-level HTTP normalization options to specify the types of encoding that are normalized for HTTP traffic, and to cause the system to generate events against traffic containing this type of encoding.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

ASCII Encoding

Decodes encoded ASCII characters and specifies whether the rules engine generates an event on ASCII-encoded URIs.

You can enable rule 119:1 to generate events for this option. See Setting Rule States, page 27-19 for more information.

UTF-8 Encoding

Decodes standard UTF-8 Unicode sequences in the URI.

You can enable rule 119:6 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Microsoft %U Encoding

Decodes the IIS %u encoding scheme that uses %u followed by four characters where the 4 characters are a hex encoded value that correlates to an IIS Unicode codepoint.

Tip Legitimate clients rarely use %u encodings, so Cisco recommends decoding HTTP traffic encoded with %u encodings.

You can enable rule 119:3 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Bare Byte UTF-8 Encoding

Decodes bare byte encoding, which uses non-ASCII characters as valid values in decoding UTF-8 values.

Tip Bare byte encoding allows the user to emulate an IIS server and interpret non-standard encodings correctly. Cisco recommends enabling this option because no legitimate clients encode UTF-8 this way.

You can enable rule 119:4 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Microsoft IIS Encoding

Decodes using Unicode codepoint mapping.

Tip Cisco recommends enabling this option, because it is seen mainly in attacks and evasion attempts.

You can enable rule 119:7 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Double Encoding

Decodes IIS double encoded traffic by making two passes through the request URI performing decodes in each one. Cisco recommends enabling this option because it is usually found only in attack scenarios.

You can enable rule 119:2 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Multi-Slash Obfuscation

Normalizes multiple slashes in a row into a single slash.

You can enable rule 119:8 to generate events for this option. See Setting Rule States, page 27-19 for more information.

IIS Backslash Obfuscation

Normalizes backslashes to forward slashes.

You can enable rule 119:9 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Directory Traversal

Normalizes directory traversals and self-referential directories. If you enable the accompanying preprocessor rules to generate events against this type of traffic, it may generate false positives because some web sites refer to files using directory traversals.

You can enable rules 119:10 and 119:11 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Tab Obfuscation

Normalizes the non-RFC standard of using a tab for a space delimiter. Apache and other non-IIS web servers use the tab character (0x09) as a delimiter in URLs.

Note Regardless of the configuration for this option, the HTTP Inspect preprocessor treats a tab as white space if a space character (0x20) precedes it.

You can enable rule 119:12 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Invalid RFC Delimiter

Normalizes line breaks (\n) in URI data.

You can enable rule 119:13 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Webroot Directory Traversal

Detects directory traversals that traverse past the initial directory in the URL.

You can enable rule 119:18 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Tab URI Delimiter

Turns on the use of the tab character (0x09) as a delimiter for a URI. Apache, newer versions of IIS, and some other web servers use the tab character as a delimiter in URLs.



Note Regardless of the configuration for this option, the HTTP Inspect preprocessor treats a tab as white space if a space character (0x20) precedes it.

Non-RFC characters

Detects the non-RFC character list you add in the corresponding field when it appears within incoming or outgoing URI data. When modifying this field, use the hexadecimal format that represents the byte character. If and when you configure this option, set the value with care. Using a character that is very common may overwhelm you with events.

You can enable rule 119:14 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Max Chunk Encoding Size

Detects abnormally large chunk sizes in URI data.

You can enable rules 119:16 and 119:22 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Disable Pipeline Decoding

Disables HTTP decoding for pipelined requests. When this option is disabled, performance is enhanced because HTTP requests waiting in the pipeline are not decoded or analyzed, and are only inspected using generic pattern matching.

Non-Strict URI Parsing

Enables non-strict URI parsing. Use this option only on servers that will accept non-standard URIs in the format "GET /index.html abc xo qr \n". Using this option, the decoder assumes that the URI is between the first and second space, even if there is no valid HTTP identifier after the second space.

Extended ASCII Encoding

Enables parsing of extended ASCII characters in an HTTP request URI. Note that this option is available in custom server profiles only, and not in the default profiles provided for Apache, IIS, or all servers.

Configuring HTTP Server Options

License: Protection

Use the following procedure to configure HTTP server options. For more information on the HTTP server options, see Selecting Server-Level HTTP Normalization Options, page 22-33 and Selecting Server-Level HTTP Normalization Encoding Options, page 22-41.

To configure server-level HTTP configuration options:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **HTTP Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The HTTP Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 You have two options:

- Add a new server profile. Click the add icon ( ) next to **Servers** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses for the client in the **Server Address** field and click **OK**.

  You can specify a single IP address or address block, or a comma-separated list of either or both. You can include up to 496 characters in a list, specify a total of 256 address entries for all server profiles, and create a total of 255 profiles including the default profile. For information on using IPv4 and IPv6 address blocks in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

  Note that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks, and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

  A new entry appears in the list of servers on the left side of the page, highlighted to indicate that it is selected, and the Configuration section updates to reflect the current configuration for the profile you added.

- Modify the settings for an existing profile. Click the configured address for a profile you have added under **Servers** on the left side of the page, or click **default**.

  Your selection is highlighted and the Configuration section updates to display the current configuration for the profile you selected. To delete an existing profile, click the delete icon ( ) next to the profile you want to remove.

Step 10 Optionally, modify the address or addresses listed in the **Networks** field and click any other area of the page.

The highlighted address updates on the left side of the page.

Note that you cannot modify the setting for **Networks** in the default profile. The default profile applies to all servers on your network that are not identified in another profile.

Step 11 In the **Ports** field, list the ports whose traffic you want to inspect with HTTP Inspect. Separate multiple ports with commas.

Step 12 You can modify any of the other options described in Selecting Server-Level HTTP Normalization Options, page 22-33.

Step 13 Select a server profile as follows:

- Select **Custom** to create your own server profile (see Selecting Server-Level HTTP Normalization Encoding Options, page 22-41 for more information).
- Select **All** to use the standard default profile, appropriate for all servers.
- Select **IIS** to use the default IIS profile.
- Select **Apache** to use the default Apache profile.

Step 14 If you selected **Custom**, the custom options appear.

Step 15 Configure the HTTP decoding options you want in your profile.

See Selecting Server-Level HTTP Normalization Options, page 22-33 for details on available normalization options.

Step 16 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Enabling Additional HTTP Inspect Preprocessor Rules

License: Protection

You can enable the rules in the **Preprocessor Rule GID:SID** column of the following table to generate events for HTTP Inspect preprocessor rules that are not associated with specific configuration options. See Setting Rule States, page 27-19 for more information.

Table 22-9 Additional HTTP Inspect Preprocessor Rules

Preprocessor Rule GID:SID	Description
120:5	Generates an event when UTF-7 encoding is encountered in HTTP response traffic; UTF-7 should only appear where 7-bit parity is required, such as in SMTP traffic.
119:21	Generates an event when an HTTP request header has more than one `content-length` field.
119:24	Generates an event when an HTTP request has more than one Host header.
119:28 120:8	When enabled, these rules do not generate events.
119:32	Generates an event when HTTP version 0.9 is encountered in traffic. Note that the TCP stream configuration must also be enabled. See Using TCP Stream Preprocessing, page 24-20.
119:33	Generates an event when an HTTP URI includes an unescaped space.
119:34	Generates an event when a TCP connection contains 24 or more pipelined HTTP requests.

Using the Sun RPC Preprocessor

License: Protection

RPC (Remote Procedure Call) normalization takes fragmented RPC records and normalizes them to a single record so the rules engine can inspect the complete record. For example, an attacker may attempt to discover the port where RPC `admind` runs. Some UNIX hosts use RPC `admind` to perform remote distributed system tasks. If the host performs weak authentication, a malicious user could take control of remote administration. The standard text rule (generator ID: 1) with the Snort ID (SID) 575 detects this attack by searching for content in specific locations to identify inappropriate `portmap GETPORT` requests.

Ports

Specify the ports whose traffic you want to normalize. In the interface, list multiple ports separated by commas. Typical RPC ports are 111 and 32771. If your network sends RPC traffic to other ports, consider adding them.

Detect fragmented RPC records

Detects RPC fragmented records.

You can enable rules 106:1 and 106:5 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect multiple records in one packet

Detects more than one RPC request per packet (or reassembled packet).

You can enable rule 106:2 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect fragmented record sums which exceed one fragment

Detects reassembled fragment record lengths that exceed the current packet length.

You can enable rule 106:3 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect single fragment records which exceed the size of one packet

Detects partial records

You can enable rule 106:4 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Configuring the Sun RPC Preprocessor

License: Protection

You can use the following procedure to configure the Sun RPC preprocessor. For more information on the Sun RPC preprocessor configuration options, see Using the Sun RPC Preprocessor, page 22-46.

To configure the Sun RPC preprocessor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **Sun RPC Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Sun RPC Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 In the **Ports** field, type the port numbers where you want to decode RPC traffic. Separate multiple ports with commas.

Step 10 You can select or clear any of the following detection options on the Sun RPC Configuration page:

- **Detect fragmented RPC records**
- **Detect multiple records in one packet**
- **Detect fragmented record sums which exceed one packet**
- **Detect single fragment records which exceed the size of one packet**

Step 11 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Decoding the Session Initiation Protocol

License: Protection

The Session Initiation Protocol (SIP) provides call setup, modification, and teardown of one or more sessions for one or more users of such client applications as Internet telephony, multimedia conferencing, instant messaging, online gaming, and file transfer. A *method* field in each SIP request identifies the purpose of the request, and a Request-URI specifies where to send the request. A status code in each SIP response indicates the outcome of the requested action.

After calls are set up using SIP, the Real-time Transport Protocol (RTP) is responsible for subsequent audio and video communication; this part of the session is sometimes referred to as the call channel, the data channel, or the audio/video data channel. RTP uses the Session Description Protocol (SDP) within the SIP message body for data-channel parameter negotiation, session announcement, and session invitation.

The SIP preprocessor is responsible for:

- decoding and analyzing SIP 2.0 traffic
- extracting the SIP header and message body, including SDP data when present, and passing the extracted data to the rules engine for further inspection
- generating events when the following conditions are detected and the corresponding preprocessor rules are enabled: anomalies and known vulnerabilities in SIP packets; out-of-order and invalid call sequences
- optionally ignoring the call channel

The preprocessor identifies the RTP channel based on the port identified in the SDP message, which is embedded in the SIP message body, but the preprocessor does not provide RTP protocol inspection.

Note the following when using the SIP preprocessor:

- UDP typically carries media sessions supported by SIP. UDP stream preprocessing provides SIP session tracking for the SIP preprocessor.
- SIP rule keywords allow you to point to the SIP packet header or message body and to limit detection to packets for specific SIP methods or status codes. For more information, see SIP Keywords, page 30-62.
- When enabled, the preprocessor generates no events before sending the extracted data to the rules engine unless you also enable the accompanying rules with generator ID (GID) 140. See Setting Rule States, page 27-19 for more information.

See the following sections for more information:

- Selecting SIP Preprocessor Options, page 22-49
- Configuring the SIP Preprocessor, page 22-50
- Enabling Additional SIP Preprocessor Rules, page 22-51

Selecting SIP Preprocessor Options

License: Protection

The following list describes SIP preprocessor options you can modify.

For the **Maximum Request URI Length**, **Maximum Call ID Length**, **Maximum Request Name Length**, **Maximum From Length**, **Maximum To Length**, **Maximum Via Length**, **Maximum Contact Length**, and **Maximum Content Length** options, you can specify from 1 to 65535 bytes, or 0 to disable event generation for the option regardless of whether the associated rule is enabled.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Ports

Specifies the ports to inspect for SIP traffic. You can specify an integer from 0 to 65535. Separate multiple port numbers with commas.

Methods to Check

Specifies SIP methods to detect. You can specify any of the following currently defined SIP methods:

```
ack, benotify, bye, cancel, do, info, invite, join, message,
notify, options, prack, publish, quath, refer, register,
service, sprack, subscribe, unsubscribe, update
```

Methods are case-insensitive. The method name can include alphabetic characters, numbers, and the underscore character. No other special characters are permitted. Separate multiple methods with commas.

Because new SIP methods might be defined in the future, your configuration can include an alphabetic string that is not currently defined. The system supports up to 32 methods, including the 21 currently defined methods and an additional 11 methods. The system ignores any undefined methods that you might configure.

Note that, in addition to any methods you specify for this option, the 32 total methods includes methods specified using the `sip_method` keyword in intrusion rules. See sip_method, page 30-62 for more information.

Maximum Dialogs within a Session

Specifies the maximum number of dialogs allowed within a stream session. If more dialogs than this number are created, the oldest dialogs are dropped until the number of dialogs does not exceed the maximum number specified; an event also triggers when rule 140:27 is enabled.

You can specify an integer from 1 to 4194303.

Maximum Request URI Length

Specifies the maximum number of bytes to allow in the Request-URI header field. A longer URI triggers an event when rule 140:3 is enabled. The request URI field indicates the destination path or page for the request.

Maximum Call ID Length

Specifies the maximum number of bytes to allow in the request or response Call-ID header field. A longer Call-ID triggers an event when rule 140:5 is enabled. The Call-ID field uniquely identifies the SIP session in requests and responses.

Maximum Request Name Length

Specifies the maximum number of bytes to allow in the request name, which is the name of the method specified in the CSeq transaction identifier. A longer request name triggers an event when rule 140:7 is enabled.

Maximum From Length

Specifies the maximum number of bytes to allow in the request or response From header field. A longer From triggers an event when rule 140:9 is enabled. The From field identifies the message initiator.

Maximum To Length

Specifies the maximum number of bytes to allow in the request or response To header field. A longer To triggers an event when rule 140:11 is enabled. The To field identifies the message recipient.

Maximum Via Length

Specifies the maximum number of bytes to allow in the request or response Via header field. A longer Via triggers an event when rule 140:13 is enabled. The Via field provides the path followed by the request and, in a response, receipt information.

Maximum Contact Length

Specifies the maximum number of bytes to allow in the request or response Contact header field. A longer Contact triggers an event when rule 140:15 is enabled. The Contact field provides a URI that specifies the location to contact with subsequent messages.

Maximum Content Length

Specifies the maximum number of bytes to allow in the content of the request or response message body. Longer content triggers an event when rule 140:16 is enabled.

Ignore Audio/Video Data Channel

Enables and disables inspection of data channel traffic. Note that the preprocessor continues inspection of other non-data-channel SIP traffic when you enable this option.

Configuring the SIP Preprocessor

License: Protection

Use the following procedure to configure the SIP preprocessor.

To configure the SIP preprocessor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **SIP Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The SIP Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 You can modify any of the options described in Selecting SIP Preprocessor Options, page 22-49.

Step 10 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Enabling Additional SIP Preprocessor Rules

License: Protection

The SIP preprocessor rules in the following table are not associated with specific configuration options. As with other SIP preprocessor rules, you must enable these rules if you want them to generate events. See Setting Rule States, page 27-19 for information on enabling rules.

Table 22-10 Additional SIP Preprocessor Rules

Preprocessor Rule GID:SID	Description
140:1	Generates an event when the preprocessor is monitoring the maximum number of SIP sessions allowed by the system.
140:2	Generates an event when the required Request_URI field is empty in a SIP request.

Table 22-10 Additional SIP Preprocessor Rules (continued)

Preprocessor Rule GID:SID	Description
140:4	Generates an event when the Call-ID header field is empty in a SIP request or response.
140:6	Generates an event when the value for the sequence number in the SIP request or response CSeq field is not a 32-bit unsigned integer less than 231.
140:8	Generates an event an event when the From header field is empty in a SIP request or response.
140:10	Generates an event when the To header field is empty in a SIP request or response.
140:12	Generates an event when the Via header field is empty in a SIP request or response
140:14	Generates an event when the required Contact header field is empty in a SIP request or response.
140:17	Generates an event when a single SIP request or response packet in UDP traffic contains multiple messages. Note that older SIP versions supported multiple messages, but SIP 2.0 supports only one message per packet.
140:18	Generates an event when the actual length of the message body in a SIP request or response in UDP traffic does not match the value specified in the Content-Length header field in a SIP request or response.
140:19	Generates an event when the preprocessor does not recognize a method name in the CSeq field of a SIP response.
140:20	Generates an event when the SIP server does not challenge an authenticated invite message. Note that this occurs in the case of the InviteReplay billing attack.
140:21	Generates an event when session information changes before the call is set up. Note that this occurs in the case of the FakeBusy billing attack.
140:22	Generates an event when the response status code is not a three-digit number.
140:23	Generates an event when the Content-Type header field does not specify a content type and the message body contains data.
140:24	Generates an event when the SIP version is not 1, 1.1, or 2.0.
140:25	Generates an event when the method specified in the CSeq header and the method field do not match in a SIP request.
140:26	Generates an event when the preprocessor does not recognize the method named in the SIP request method field.

Configuring the GTP Command Channel

License: Protection

The General Service Packet Radio (GPRS) Tunneling Protocol (GTP) provides communication over a GTP core network. The GTP preprocessor detects anomalies in GTP traffic and forwards command channel signaling messages to the rules engine for inspection. You can use the `gtp_version`, `gtp_type`, and `gtp_info` rule keywords to inspect GTP command channel traffic for exploits.

A single configuration option allows you to modify the default setting for the ports that the preprocessor inspects for GTP command channel messages.

You must enable the GTP preprocessor rules in the following table if you want them to generate events. See Setting Rule States, page 27-19 for information on enabling rules.

Table 22-11 GTP Preprocessor Rules

Preprocessor Rule GID:SID	Description
143:1	Generates an event when the preprocessor detects an invalid message length.
143:2	Generates an event when the preprocessor detects an invalid information element length.
143:3	Generates an event when the preprocessor detects information elements that are out of order.

You can use the following procedure to modify the ports the GTP preprocessor monitors for GTP command messages.

To configure the GTP command channel:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **GTP Command Channel Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The GTP Command Channel Configuration page appears.

Step 9 Optionally, modify the ports that the preprocessor inspects for GTP command messages. You can specify an integer from 0 to 65535. Use commas to separate multiple ports.

Step 10 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Decoding IMAP Traffic

License: Protection

The Internet Message Application Protocol (IMAP) is used to retrieve email from a remote IMAP server. The IMAP preprocessor inspects server-to-client IMAP4 traffic and, when associated preprocessor rules are enabled, generates events on anomalous traffic. The preprocessor can also extract and decode email attachments in client-to-server IMAP4 traffic and send the attachment data to the rules engine. You can use the `file_data` keyword in an intrusion rule to point to the attachment data. See Pointing to a Specific Payload Type, page 30-98 for more information.

Extraction and decoding include multiple attachments, when present, and large attachments that span multiple packets.

If you want IMAP preprocessor rules to generate events, you must enable the rules. IMAP preprocessor rules have a generator ID (GID) of 141.See Setting Rule States, page 27-19 for more information.

See the following sections for more information:

- Selecting IMAP Preprocessor Options, page 22-54
- Configuring the IMAP Preprocessor, page 22-55
- Enabling Additional IMAP Preprocessor Rules, page 22-56

Selecting IMAP Preprocessor Options

License: Protection

The following list describes the IMAP preprocessor options you can modify.

Note that decoding, or extraction when the MIME email attachment does not require decoding, includes multiple attachments when present, and large attachments that span multiple packets.

Note also that the highest value is used when the values for the **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** options are different in:

- the default network analysis policy
- any other custom network analysis policy invoked by network analysis rules in the same access control policy

See Setting the Default Network Analysis Policy for Access Control, page 20-3 and Specifying Traffic to Preprocess Using Network Analysis Rules, page 20-4 for more information.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Ports

Specifies the ports to inspect for IMAP traffic. You can specify an integer from 0 to 65535. Separate multiple port numbers with commas.

Base64 Decoding Depth

Specifies the maximum number of bytes to extract and decode from each Base64 encoded MIME email attachment. You can specify from 1 to 65535 bytes, or specify 0 to decode all the Base64 data. Specify -1 to ignore Base64 data.

Note that positive values not divisible by 4 are rounded up to the next multiple of 4 except for the values 65533, 65534, and 65535, which are rounded down to 65532.

When Base64 decoding is enabled, you can enable rule 141:4 to generate an event when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

7-Bit/8-Bit/Binary Decoding Depth

Specifies the maximum bytes of data to extract from each MIME email attachment that does not require decoding. These attachment types include 7-bit, 8-bit, binary, and various multipart content types such as plain text, jpeg images, mp3 files, and so on. You can specify from 1 to 65535 bytes, or specify 0 to extract all data in the packet. Specify -1 to ignore non-decoded data.

Quoted-Printable Decoding Depth

Specifies the maximum number of bytes to extract and decode from each quoted-printable (QP) encoded MIME email attachment. You can specify from 1 to 65535 bytes, or specify 0 to decode all QP encoded data in the packet. Specify -1 to ignore QP encoded data.

When quoted-printable decoding is enabled, you can enable rule 141:6 to generate an event when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

Unix-to-Unix Decoding Depth

Specifies the maximum number of bytes to extract and decode from each Unix-to-Unix encoded (uuencoded) email attachment. You can specify from 1 to 65535 bytes, or specify 0 to decode all uuencoded data in the packet. Specify -1 to ignore uuencoded data.

When Unix-to-Unix decoding is enabled, you can enable rule 141:7 to generate an event when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data.

Configuring the IMAP Preprocessor

License: Protection

Use the following procedure to configure the IMAP preprocessor. For additional information on IMAP preprocessor configuration options, see Selecting IMAP Preprocessor Options, page 22-54.

To configure the IMAP preprocessor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **IMAP Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The IMAP Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 Specify the **Ports** where IMAP traffic should be decoded. Separate multiple port numbers with commas.

Step 10 Specify the maximum bytes of data to extract and decode from any combination of the following email attachment types:

- **Base64 Decoding Depth**
- **7-Bit/8-Bit/Binary Decoding Depth** (includes various multipart content types such as plain text, jpeg images, mp3 files, and so on)
- **Quoted-Printable Decoding Depth**
- **Unix-to-Unix Decoding Depth**

For each type, you can specify from 1 to 65535 bytes, or specify 0 to extract and, when necessary, decode all data in the packet. Specify -1 to ignore data for an attachment type.

You can use the `file_data` rule keyword in intrusion rules to inspect the attachment data. See Pointing to a Specific Payload Type, page 30-98 for more information.

Step 11 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Enabling Additional IMAP Preprocessor Rules

License: Protection

The IMAP preprocessor rules in the following table are not associated with specific configuration options. As with other IMAP preprocessor rules, you must enable these rules if you want them to generate events. See Setting Rule States, page 27-19 for information on enabling rules.

Table 22-12 Additional IMAP Preprocessor Rules

Preprocessor Rule GID:SID	Description
141:1	Generates an event when the preprocessor detects a client command that is not defined in RFC 3501.
141:2	Generates an event when the preprocessor detects a server response that is not defined in RFC 3501.
141:3	Generates an event when the preprocessor is using the maximum amount of memory allowed by the system. At this point, the preprocessor stops decoding until memory becomes available.

Decoding POP Traffic

License: Protection

The Post Office Protocol (POP) is used to retrieve email from a remote POP mail server. The POP preprocessor inspects server-to-client POP3 traffic and, when associated preprocessor rules are enabled, generates events on anomalous traffic. The preprocessor can also extract and decode email attachments in client-to-server POP3 traffic and send the attachment data to the rules engine. You can use the `file_data` keyword in an intrusion rule to point to attachment data. See Pointing to a Specific Payload Type, page 30-98 for more information.

Extraction and decoding include multiple attachments, when present, and large attachments that span multiple packets.

If you want POP preprocessor rules to generate events, you must enable the rules. POP preprocessor rules have a generator ID (GID) of 142. See Setting Rule States, page 27-19 for more information.

See the following sections for more information:

- Selecting POP Preprocessor Options, page 22-57
- Configuring the POP Preprocessor, page 22-58
- Enabling Additional POP Preprocessor Rules, page 22-59

Selecting POP Preprocessor Options

License: Protection

The following list describes the POP preprocessor options you can modify.

Note that decoding, or extraction when the MIME email attachment does not require decoding, includes multiple attachments when present, and large attachments that span multiple packets.

Note also that when the values for the **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** options are different in an intrusion policy associated with the default action of an access control policy and intrusion policies associated with access control rules, the highest value is used.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Ports

Specifies the ports to inspect for POP traffic. You can specify an integer from 0 to 65535. Separate multiple port numbers with commas.

Base64 Decoding Depth

Specifies the maximum number of bytes to extract and decode from each Base64 encoded MIME email attachment. You can specify from 1 to 65535 bytes, or specify 0 to decode all the Base64 data. Specify -1 to ignore Base64 data.

Note that positive values not divisible by 4 are rounded up to the next multiple of 4 except for the values 65533, 65534, and 65535, which are rounded down to 65532.

When Base64 decoding is enabled, you can enable rule 142:4 to generate an event when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data. See Setting Rule States, page 27-19 for more information.

7-Bit/8-Bit/Binary Decoding Depth

Specifies the maximum bytes of data to extract from each MIME email attachment that does not require decoding. These attachment types include 7-bit, 8-bit, binary, and various multipart content types such as plain text, jpeg images, mp3 files, and so on. You can specify from 1 to 65535 bytes, or specify 0 to extract all data in the packet. Specify -1 to ignore non-decoded data.

Quoted-Printable Decoding Depth

Specifies the maximum number of bytes to extract and decode from each quoted-printable (QP) encoded MIME email attachment. You can specify from 1 to 65535 bytes, or specify 0 to decode all QP encoded data in the packet. Specify -1 to ignore QP encoded data.

When quoted-printable decoding is enabled, you can enable rule 142:6 to generate an event when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data. See Setting Rule States, page 27-19 for more information.

Unix-to-Unix Decoding Depth

Specifies the maximum number of bytes to extract and decode from each Unix-to-Unix encoded (uuencoded) email attachment. You can specify from 1 to 65535 bytes, or specify 0 to decode all uuencoded data in the packet. Specify -1 to ignore uuencoded data.

When Unix-to-Unix decoding is enabled, you can enable rule 142:7 to generate an event when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data. See Setting Rule States, page 27-19 for more information.

Configuring the POP Preprocessor

License: Protection

Use the following procedure to configure the POP preprocessor. For additional information on POP preprocessor configuration options, see Selecting POP Preprocessor Options, page 22-57.

To configure the POP preprocessor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **POP Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The POP Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 Specify the **Ports** where IMAP traffic should be decoded. Separate multiple port numbers with commas.

Step 10 Specify the maximum bytes of data to extract and decode from any combination of the following email attachment types:

- **Base64 Decoding Depth**
- **7-Bit/8-Bit/Binary Decoding Depth** (includes various multipart content types such as plain text, jpeg images, mp3 files, and so on)
- **Quoted-Printable Decoding Depth**
- **Unix-to-Unix Decoding Depth**

For each type, you can specify from 1 to 65535 bytes, or specify 0 to extract and, when necessary, decode all data in the packet. Specify -1 to ignore data for an attachment type.

You can use the `file_data` rule keyword in intrusion rules to inspect the attachment data. See Pointing to a Specific Payload Type, page 30-98 for more information.

Step 11 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Enabling Additional POP Preprocessor Rules

License: Protection

The POP preprocessor rules in the following table are not associated with specific configuration options. As with other POP preprocessor rules, you must enable these rules if you want them to generate events. See Setting Rule States, page 27-19 for information on enabling rules.

Table 22-13 Additional POP Preprocessor Rules

Preprocessor Rule GID:SID	Description
142:1	Generates an event when the preprocessor detects a client command that is not defined in RFC 1939.
142:2	Generates an event when the preprocessor detects a server response that is not defined in RFC 1939.
142:3	Generates an event when the preprocessor is using the maximum amount of memory allowed by the system. At this point, the preprocessor stops decoding until memory becomes available.

Decoding SMTP Traffic

License: Protection

The SMTP preprocessor instructs the rules engine to normalize SMTP commands. The preprocessor can also extract and decode email attachments in client-to-server traffic and, depending on the software version, extract email file names, addresses, and header data to provide context when displaying intrusion events triggered by SMTP traffic.

Note the following when using the SMTP preprocessor:

- You must enable SMTP preprocessor rules, which have a generator ID (GID) of 124, if you want these rules to generate events. See Setting Rule States, page 27-19 for more information.

For more information, see the following sections:

- Understanding SMTP Decoding, page 22-60
- Configuring SMTP Decoding, page 22-64
- Enabling SMTP Maximum Decoding Memory Alerting, page 22-67

Understanding SMTP Decoding

License: Protection

You can enable or disable normalization, and you can configure options to control the types of anomalous traffic the SMTP decoder detects.

Note that decoding, or extraction when the MIME email attachment does not require decoding, includes multiple attachments when present, and large attachments that span multiple packets.

Note also that when the values for the **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth** options are different in an intrusion policy associated with the default action of an access control policy and intrusion policies associated with access control rules, the highest value is used.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Ports

Specifies the ports whose SMTP traffic you want to normalize. You can specify an integer from 0 to 65535. Separate multiple ports with commas.

Stateful Inspection

When selected, causes SMTP decoder to save state and provide session context for individual packets and only inspects reassembled sessions. When cleared, analyzes each individual packet without session context.

Normalize

When set to All, normalizes all commands. Checks for more than one space character after a command.

When set to None, normalizes no commands.

When set to Cmds, normalizes the commands listed in **Custom Commands**.

Custom Commands

When **Normalize** is set to Cmds, normalizes the listed commands.

Specify commands which should be normalized in the text box. Checks for more than one space character after a command.

The space (ASCII 0x20) and tab (ASCII 0x09) characters count as space characters for normalization purposes.

Ignore Data

Does not process mail data; processes only MIME mail header data.

Ignore TLS Data

Does not process data encrypted under the Transport Layer Security protocol.

No Alerts

Disables intrusion events when accompanying preprocessor rules are enabled.

Detect Unknown Commands

Detects unknown commands in SMTP traffic.

You can enable rules 124:5 and 124:6 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Max Command Line Len

Detects when an SMTP command line is longer than this value. Specify `0` to never detect command line length.

RFC 2821, the Network Working Group specification on the Simple Mail Transfer Protocol, recommends 512 as a maximum command line length.

You can enable rule 124:1 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Max Header Line Len

Detects when an SMTP data header line is longer than this value. Specify `0` to never detect data header line length.

You can enable rules 124:2 and 124:7 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Max Response Line Len

Detects when an SMTP response line is longer than this value. Specify `0` to never detect response line length.

RFC 2821 recommends 512 as a maximum response line length.

You can enable rule 124:3 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Alt Max Command Line Len

Detects when the SMTP command line for any of the specified commands is longer than this value. Specify `0` to never detect command line length for the specified commands. Different default line lengths are set for numerous commands.

This setting overrides the Max Command Line Len setting for the specified commands.

You can enable rule 124:3 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Invalid Commands

Detects if these commands are sent from the client side.

You can enable rule 124:5 and 124:6 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Valid Commands

Permits commands in this list.

Even if this list is empty, the preprocessor permits the following valid commands: ATRN AUTH BDAT DATA DEBUG EHLO EMAL ESAM ESND ESOM ETRN EVFY EXPN HELO HELP IDENT MAIL NOOP ONEX QUEU QUIT RCPT RSET SAML SEND SIZE SOML STARTTLS TICK TIME TURN TURNME VERB VRFY XADR XAUTH XCIR XEXCH50 X-EXPS XGEN XLICENSE X-LINK2STATE XQUE XSTA XTRN XUSR



Note RCPT TO and MAIL FROM are SMTP commands. The preprocessor configuration uses command names of RCPT and MAIL, respectively. Within the code, the preprocessor maps RCPT and MAIL to the correct command name.

You can enable rule 124:4 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Data Commands

Lists commands that initiate sending data in the same way the SMTP DATA command sends data per RFC 5321. Separate multiple commands with spaces.

Binary Data Commands

Lists commands that initiate sending data in a way that is similar to how the BDAT command sends data per RFC 3030. Separate multiple commands with spaces.

Authentication Commands

Lists commands that initiate an authentication exchange between client and server. Separate multiple commands with spaces.

Detect xlink2state

Detects packets that are part of X-Link2State Microsoft Exchange buffer data overflow attacks. In inline deployments, the system can also drop those packets.

You can enable rule 124:8 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Base64 Decoding Depth

When **Ignore Data** is disabled, specifies the maximum number of bytes to extract and decode from each Base64 encoded MIME email attachment. You can specify from 1 to 65535 bytes, or specify 0 to decode all the Base64 data. Specify -1 to ignore Base64 data. The preprocessor will not decode data when **Ignore Data** is selected.

Note that positive values not divisible by 4 are rounded up to the next multiple of 4 except for the values 65533, 65534, and 65535, which are rounded down to 65532.

When Base64 decoding is enabled, you can enable rule 124:10 to generate an event when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data. See Setting Rule States, page 27-19 for more information.

Note that this option replaces the deprecated options **Enable MIME Decoding** and **Maximum MIME Decoding Depth**, which are still supported in existing intrusion policies for backward compatibility.

7-Bit/8-Bit/Binary Decoding Depth

When **Ignore Data** is disabled, specifies the maximum bytes of data to extract from each MIME email attachment that does not require decoding. These attachment types include 7-bit, 8-bit, binary, and various multipart content types such as plain text, jpeg images, mp3 files, and so on. You can specify from 1 to 65535 bytes, or specify 0 to extract all data in the packet. Specify -1 to ignore non-decoded data. The preprocessor will not extract data when **Ignore Data** is selected.

Quoted-Printable Decoding Depth

When **Ignore Data** is disabled, specifies the maximum number of bytes to extract and decode from each quoted-printable (QP) encoded MIME email attachment.

You can specify from 1 to 65535 bytes, or specify 0 to decode all QP encoded data in the packet. Specify -1 to ignore QP encoded data. The preprocessor will not decode data when **Ignore Data** is selected.

When quoted-printable decoding is enabled, you can enable rule 124:11 to generate an event when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data. See Setting Rule States, page 27-19 for more information.

Unix-to-Unix Decoding Depth

When **Ignore Data** is disabled, specifies the maximum number of bytes to extract and decode from each Unix-to-Unix encoded (uuencoded) email attachment. You can specify from 1 to 65535 bytes, or specify 0 to decode all uuencoded data in the packet. Specify -1 to ignore uuencoded data. The preprocessor will not decode data when **Ignore Data** is selected.

When Unix-to-Unix decoding is enabled, you can enable rule 124:13 to generate an event when decoding fails; decoding could fail, for example, because of incorrect encoding or corrupted data. See Setting Rule States, page 27-19 for more information.

Log MIME Attachment Names

Enables extraction of MIME attachment file names from the MIME Content-Disposition header and associates the file names with all intrusion events generated for the session. Multiple file names are supported.

When this option is enabled, you can view file names associated with events in the Email Attachment column of the intrusion events table view. See Viewing Events, page 37-1 for more information.

Log To Addresses

Enables extraction of recipient email addresses from the SMTP RCPT TO command and associates the recipient addresses with all intrusion events generated for the session. Multiple recipients are supported.

When this option is enabled, you can view recipients associated with events in the Email Recipient column of the intrusion events table view. See Viewing Events, page 37-1 for more information.

Log From Addresses

Enables extraction of sender email addresses from the SMTP MAIL FROM command and associates the sender addresses with all intrusion events generated for the session. Multiple sender addresses are supported.

When this option is enabled, you can view senders associated with events in the Email Sender column of the intrusion events table view. See Viewing Events, page 37-1 for more information.

Log Headers

Enables extraction of email headers. The number of bytes to extract is determined by the value specified for **Header Log Depth**.

You can use the `content` or `protected_content` keyword to write intrusion rules that use email header data as a pattern. You can also view the extracted email header in the intrusion event packet view. See Viewing Events, page 37-1 for more information.

Header Log Depth

Specifies the number of bytes of the email header to extract when **Log Headers** is enabled. You can specify 0 to 20480 bytes. A value of 0 disables **Log Headers**.

Configuring SMTP Decoding

License: Protection

You can use the SMTP Configuration page of an intrusion policy to configure SMTP normalization. For more information on SMTP preprocessor configuration options, see Understanding SMTP Decoding, page 22-60.

To configure SMTP decoding options:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **SMTP Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The SMTP Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 Specify the **Ports** where SMTP traffic should be decoded, separated by commas.

Step 10 Select **Stateful Inspection** to examine reassembled TCP streams containing SMTP packets. Clear **Stateful Inspection** to inspect only unreassembled SMTP packets.

Step 11 Configure the normalization options:

- To normalize all commands, select **All**.
- To normalize only commands specified by **Custom Commands**, select **Cmds** and specify the commands to normalize. Separate commands with spaces.
- To normalize no commands, select **None**.
- To ignore mail data except for MIME mail header data, check **Ignore Data**.
- To ignore data encrypted under the Transport Security Layer protocol, check **Ignore TLS Data**.
- To disable generating events when accompanying preprocessor rules are enabled, check **No Alerts**.
- To detect unknown commands in SMTP data, select **Detect Unknown Commands**.

Step 12 Specify a maximum command line length in the **Max Command Line Len** field.

Step 13 Specify a maximum data header line length in the **Max Header Line Len** field.

Step 14 Specify a maximum response line length in the **Max Response Line Len** field.

Note RCPT TO and MAIL FROM are SMTP commands. The preprocessor configuration uses command names of RCPT and MAIL, respectively. Within the code, the preprocessor maps RCPT and MAIL to the correct command name.

Step 15 If needed, click **Add** next to **Alt Max Command Line Len** to add commands where you want to specify an alternate maximum command line length, then specify the line length and the command or commands, separated by spaces, where you want that length to be enforced.

Step 16 Specify any commands that you want to treat as invalid and detect in the **Invalid Commands** field. Separate commands with spaces.

Step 17 Specify any commands that you want to treat as valid in the **Valid Commands** field. Separate commands with spaces.

Note Even if the **Valid Commands** list is empty, the preprocessor treats the following commands as valid: ATRN, AUTH, BDAT, DATA, DEBUG, EHLO, EMAL, ESAM, ESND, ESOM, ETRN, EVFY, EXPN, HELO, HELP, IDENT, MAIL, NOOP, QUIT, RCPT, RSET, SAML, SOML, SEND, ONEX, QUEU, STARTTLS, TICK, TIME, TURN, TURNME, VERB, VRFY, X-EXPS, X-LINK2STATE, XADR, XAUTH, XCIR, XEXCH50, XGEN, XLICENSE, XQUE, XSTA, XTRN, or XUSR.

Step 18 Specify any commands that you want to initiate sending data in the same way the SMTP DATA command sends data per RFC 5321 in the **Data Commands** field. Separate commands with spaces.

Step 19 Specify any commands that initiate sending data in a way that is similar to how the BDAT command sends data per RFC 3030 in the **Binary Data Commands** field. Separate commands with spaces.

Step 20 Specify any commands that initiate an authentication exchange between client and server in the **Authentication Commands** field. Separate commands with spaces.

Step 21 To detect packets that are part of X-Link2State Microsoft Exchange buffer data overflow attacks, select **Detect xlink2state**.

Step 22 To specify the maximum bytes of data to extract and decode for different types of email attachment, specify a value for any of the following attachment types:

- **Base64 Decoding Depth**
- **7-Bit/8-Bit/Binary Decoding Depth** (includes various multipart content types such as plain text, jpeg images, mp3 files, and so on)
- **Quoted-Printable Decoding Depth**
- **Unix-to-Unix Decoding Depth**

You can specify from 1 to 65535 bytes, or specify 0 to extract and, when necessary, decode all data in the packet for that type. Specify -1 to ignore data for an attachment type.

You can use the `file_data` rule keyword in intrusion rules to inspect extracted data. See Pointing to a Specific Payload Type, page 30-98 for more information.

You must also select the SMTP **Stateful Inspection** option to extract and decode cross-packet data or data crossing multiple TCP segments.

Step 23 Configure options for associating contextual information with intrusion events triggered by SMTP traffic:

- To enable extraction of MIME attachment file names to associate with intrusion events, select **Log MIME Attachment Names**.
- To enable extraction of recipient email addresses, select **Log To Addresses**.
- To enable extraction of sender email addresses to associate with intrusion events, select **Log From Addresses**.
- To enable extraction of email headers to associate with intrusion events and for writing rules that inspect email headers, select **Log Headers**.

Note that header information is displayed in the intrusion event packet view. Note also that you can also write intrusion rules that use the `content` or `protected_content` keyword with email header data as a pattern. See Viewing Events, page 37-1 for more information.

Optionally, you can specify a **Header Log Depth** of 0 to 20480 bytes of the email header to extract. A value of 0 disables **Log Headers**.

Step 24 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Enabling SMTP Maximum Decoding Memory Alerting

License: Protection

You can enable SMTP preprocessor rule 124:9 to generate an event when the enabled preprocessor is using the maximum amount of memory allowed by the system for decoding the following types of encoded data:

- Base64
- 7-bit/8-bit/binary
- Quoted-printable
- Unix-to-Unix

When the maximum decoding memory is exceeded, the preprocessor stops decoding these types of encoded data until memory becomes available. This preprocessor rule is not associated with a single, specific configuration option. See Setting Rule States, page 27-19 for information on enabling rules.

Detecting Exploits Using the SSH Preprocessor

License: Protection

The SSH preprocessor detects the Challenge-Response Buffer Overflow exploit, the CRC-32 exploit, the SecureCRT SSH Client Buffer Overflow exploit, protocol mismatches, and incorrect SSH message direction. The preprocessor also detects any version string other than version 1 or 2.

Challenge-Response Buffer Overflow and CRC-32 attacks occur after the key exchange and are, therefore, encrypted. Both attacks send an uncharacteristically large payload of more than 20 KBytes to the server immediately after the authentication challenge. CRC-32 attacks apply only to SSH Version 1; Challenge-Response Buffer Overflow exploits apply only to SSH Version 2. The version string is read at the beginning of the session. Except for the difference in the version string, both attacks are handled in the same way.

The SecureCRT SSH exploit and protocol mismatch attacks occur when attempting to secure a connection, before the key exchange. The SecureCRT exploit sends an overly long protocol identifier string to the client that causes a buffer overflow. A protocol mismatch occurs when either a non-SSH client application attempts to connect to a secure SSH server or the server and client version numbers do not match.

You can configure the preprocessor to inspect traffic on a specified port or list of ports, or to automatically detect SSH traffic. It will continue to inspect SSH traffic until either a specified number of encrypted packets has passed within a specified number of bytes, or until a specified maximum

number of bytes is exceeded within the specified number of packets. If the maximum number of bytes is exceeded, it is assumed that a CRC-32 (SSH Version 1) or a Challenge-Response Buffer Overflow (SSH Version 2) attack has occurred. Additionally, you can detect the SecureCRT exploit, protocol mismatches, and bad message direction. Note that the preprocessor detects without configuration any version string value other than version 1 or 2.

Note the following when using the SSH preprocessor:

- You must enable SSH preprocessor rules, which have a generator ID (GID) of 128, if you want these rules to generate events. See Setting Rule States, page 27-19 for more information.
- The SSH preprocessor does not handle brute force attacks. For information on brute force attempts, see Adding Dynamic Rule States, page 27-28.

See the following sections for more information:

- Selecting SSH Preprocessor Options, page 22-68
- Configuring the SSH Preprocessor, page 22-70

Selecting SSH Preprocessor Options

License: Protection

This section describes the options you can use to configure the SSH preprocessor.

The preprocessor stops inspecting traffic for a session when either of the following occurs:

- a valid exchange between the server and the client has occurred for this number of encrypted packets; the connection continues.
- the **Number of Bytes Sent Without Server Response** is reached before the number of encrypted packets to inspect is reached; the assumption is made that there is an attack.

Each valid server response during **Number of Encrypted Packets to Inspect** resets the **Number of Bytes Sent Without Server Response** and the packet count continues.

Consider the following example SSH preprocessor configuration:

- **Server Ports**: 22
- **Autodetect Ports**: off
- **Maximum Length of Protocol Version String**: 80
- **Number of Encrypted Packets to Inspect**: 25
- **Number of Bytes Sent Without Server Response**: 19,600
- All detect options are enabled.

In the example, the preprocessor inspects traffic only on port 22. That is, auto-detection is disabled, so it inspects only on the specified port.

Additionally, the preprocessor in the example stops inspecting traffic when either of the following occurs:

- The client sends 25 encrypted packets which contain no more than 19,600 bytes, cumulative. The assumption is there is no attack.
- The client sends more than 19,600 bytes within 25 encrypted packets. In this case, the preprocessor considers the attack to be the Challenge-Response Buffer Overflow exploit because the session in the example is an SSH Version 2 session.

The preprocessor in the example will also detect any of the following that occur while it is processing traffic:

- a server overflow, triggered by a version string greater than 80 bytes and indicating a SecureCRT exploit
- a protocol mismatch
- a packet flowing in the wrong direction

Finally, the preprocessor will automatically detect any version string other than version 1 or version 2.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Server Ports

Specifies on which ports the SSH preprocessor should inspect traffic.

You can configure a single port or a comma-separated list of ports.

Autodetect Ports

Sets the preprocessor to automatically detect SSH traffic.

When this option is selected, the preprocessor inspects all traffic for an SSH version number. It stops processing when neither the client nor the server packet contains a version number. When disabled, the preprocessor inspects only the traffic identified by the **Server Ports** option.

Number of Encrypted Packets to Inspect

Specifies the number of encrypted packets to examine per session.

Setting this option to zero will allow all traffic to pass.

Reducing the number of encrypted packets to inspect may result in some attacks escaping detection. Raising the number of encrypted packets to inspect may negatively affect performance.

Number of Bytes Sent Without Server Response

Specifies the maximum number of bytes an SSH client may send to a server without getting a response before assuming there is a Challenge-Response Buffer Overflow or CRC-32 attack.

Increase the value for this option if the preprocessor generates false positives on the Challenge-Response Buffer Overflow or CRC-32 exploit.

Maximum Length of Protocol Version String

Specifies the maximum number of bytes allowed in the server's version string before considering it to be a SecureCRT exploit.

Detect Challenge-Response Buffer Overflow Attack

Enables or disables detecting the Challenge-Response Buffer Overflow exploit.

You can enable rule 128:1 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect SSH1 CRC-32 Attack

Enables or disables detecting the CRC-32 exploit.

You can enable rule 128:2 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect Server Overflow

Enables or disables detecting the SecureCRT SSH Client Buffer Overflow exploit.

You can enable rule 128:3 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect Protocol Mismatch

Enables or disables detecting protocol mismatches.

You can enable rule 128:4 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect Bad Message Direction

Enables or disables detecting when traffic flows in the wrong direction (that is, if the presumed server generates client traffic, or if a client generates server traffic).

You can enable rule 128:5 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect Payload Size Incorrect for the Given Payload

Enables or disables detecting packets with an incorrect payload size such as when the length specified in the SSH packet is not consistent with the total length specified in the IP header or the message is truncated, that is, there is not enough data for a full SSH header.

You can enable rule 128:6 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect Bad Version String

Note that, when enabled, the preprocessor detects without configuration any version string other than version 1 or 2.

You can enable rule 128:7 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Configuring the SSH Preprocessor

License: Protection

This section explains how to configure the SSH preprocessor.

To configure the SSH preprocessor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue.See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **SSH Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The SSH Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 You can modify any of the options on the SSH Configuration preprocessor page. See Selecting SSH Preprocessor Options, page 22-68 for more information.

Step 10 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Using the SSL Preprocessor

License: Protection

Using the SSL preprocessor, however, the system can analyze the contents of the handshake and key exchange messages exchanged at the beginning of an SSL session to determine when the session becomes encrypted. When SSL preprocessing is active, you can cause the system to suspend inspection of a session as soon as it becomes encrypted. You must ensure that TCP stream preprocessing is enabled to use the SSL preprocessor.

For more information, see the following sections:

- Understanding SSL Preprocessing, page 22-71
- Enabling SSL Preprocessor Rules, page 22-72
- Configuring the SSL Preprocessor, page 22-73

Understanding SSL Preprocessing

License: Protection

The SSL preprocessor stops inspection of encrypted data, which can help to eliminate false positives. The SSL preprocessor maintains state information as it inspects the SSL handshake, tracking both the state and SSL version for that session. When the preprocessor detects that a session state is encrypted, the system marks the traffic in that session as encrypted. You can configure the system to stop processing on all packets in an encrypted session when encryption is established.

For each packet, the SSL preprocessor verifies that the traffic contains an IP header, a TCP header, and a TCP payload, and that it occurs on the ports specified for SSL preprocessing. For qualifying traffic, the following scenarios determine whether the traffic is encrypted:

- the system observes all packets in a session, **Server side data is trusted** is not enabled, and the session includes a Finished message from both the server and the client and at least one packet from each side with an Application record and without an Alert record
- the system misses some of the traffic, **Server side data is trusted** is not enabled, and the session includes at least one packet from each side with an Application record that is not answered with an Alert record
- the system observes all packets in a session, **Server side data is trusted** is enabled, and the session includes a Finished message from the client and at least one packet from the client with an Application record and without an Alert record
- the system misses some of the traffic, **Server side data is trusted** is enabled, and the session includes at least one packet from the client with an Application record that is not answered with an Alert record

If you choose to stop processing on encrypted traffic, the system ignores future packets in a session after it marks the session as encrypted.

Note You can add the `ssl_state` and `ssl_version` keywords to a rule to use SSL state or version information within the rule. For more information, see Extracting SSL Information from a Session, page 30-54.

Enabling SSL Preprocessor Rules

License: Protection

When enabled, the SSL preprocessor inspects the contents of the handshake and key exchange messages exchanged at the beginning of an SSL session.

Note that you must enable SSL preprocessor rules, which have a generator ID (GID) of 137, if you want these rules to generate events. See Setting Rule States, page 27-19 for more information.

The following table describes the SSL preprocessor rules you can enable.

Table 22-14 SSL Preprocessor Rules

Preprocessor Rule GID:SID	Description
137:1	Detects a client hello after a server hello, which is invalid and considered to be anomalous behavior.
137:2	Detects a server hello without a client hello when **Server side data is trusted** is disabled, which is invalid and considered to be anomalous behavior. See Configuring the SSL Preprocessor, page 22-73 for more information.

Configuring the SSL Preprocessor

License: Protection

By default, the system attempts to inspect encrypted traffic. When you enable the SSL preprocessor, it detects when a session becomes encrypted. After the SSL preprocessor is enabled, the rules engine can invoke the preprocessor to obtain SSL state and version information. If you enable rules using the `ssl_state` and `ssl_version` keywords in an intrusion policy, you should also enable the SSL preprocessor in that policy.

In addition, you can enable the **Stop inspecting encrypted traffic** option to disable inspection and reassembly for encrypted sessions. The SSL preprocessor maintains state for the session so it can disable inspection of all traffic in the session. The system only stops inspecting traffic in encrypted sessions if SSL preprocessing is enabled **and** the **Stop inspecting encrypted traffic** option is selected.

To base identification of encrypted traffic only on server traffic, you can enable the **Server side data is trusted** option; that is, server side data is trusted to indicate that the traffic is encrypted. The SSL preprocessor typically checks both client traffic and the server responses to that traffic to determine if a session is encrypted. However, because the system may not mark a transaction as encrypted if it cannot detect both sides of a session, you can rely on the SSL server to indicate a session is encrypted. Note that when you enable the **Server side data is trusted** option you must also enable the **Stop inspecting encrypted traffic** option so the system does not continue inspecting traffic in the encrypted session.

You can specify the ports where the preprocessor monitors traffic for encrypted sessions.

Note If the SSL preprocessor detects non-SSL traffic over the ports specified for SSL monitoring, it tries to decode the traffic as SSL traffic, and then flags it as corrupt.

To configure the SSL preprocessor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **SSL Configuration** under Application Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The SSL Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 Type the ports, separated by commas, where the SSL preprocessor should monitor traffic for encrypted sessions. Only ports included in the **Ports** field will be checked for encrypted traffic.

Step 10 Click the **Stop inspecting encrypted traffic** check box to enable or disable inspection of traffic in a session after the session is marked as encrypted.

Step 11 Click the **Server side data is trusted** check box to enable or disable identification of encrypted traffic based only on the client-side traffic.

Step 12 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.



Configuring SCADA Preprocessing

You configure Supervisory Control and Data Acquisition (SCADA) preprocessors in a network analysis policy, which prepares traffic for inspection using the rules enabled in an intrusion policy. See Understanding Network Analysis and Intrusion Policies, page 18-1 for more information.

SCADA protocols monitor, control, and acquire data from industrial, infrastructure, and facility processes such as manufacturing, production, water treatment, electric power distribution, airport and shipping systems, and so on. The ASA FirePOWER module provides preprocessors for the Modbus and DNP3 SCADA protocols that you can configure as part of your network analysis policy.

If you enable a rule containing Modbus or DNP3 keywords in the corresponding intrusion policy, the system automatically uses the Modbus or DNP3 processor, respectively, with its current settings, although the preprocessor remains disabled in the network analysis policy module interface. For more information, see Modbus Keywords, page 30-74 and DNP3 Keywords, page 30-76.

See the following sections for more information:

- Configuring the Modbus Preprocessor, page 23-1
- Configuring the DNP3 Preprocessor, page 23-3

Configuring the Modbus Preprocessor

License: Protection

The Modbus protocol, which was first published in 1979 by Modicon, is a widely used SCADA protocol. The Modbus preprocessor detects anomalies in Modbus traffic and decodes the Modbus protocol for processing by the rules engine, which uses Modbus keywords to access certain protocol fields. See Modbus Keywords, page 30-74 for more information.

A single configuration option allows you to modify the default setting for the port that the preprocessor inspects for Modbus traffic.

You must enable the Modbus preprocessor rules in the following table if you want these rules to generate events. See Setting Rule States, page 27-19 for information on enabling rules.

Table 23-1 Modbus Preprocessor Rules

Preprocessor Rule GID:SID	Description
144:1	Generates an event when the length in the Modbus header does not match the length required by the Modbus function code. Each Modbus function has an expected format for requests and responses. If the length of the message does not match the expected format, this event is generated.
144:2	Generates an event when the Modbus protocol ID is non-zero. The protocol ID field is used for multiplexing other protocols with Modbus. Because the preprocessor does not process these other protocols, this event is generated instead.
144:3	Generates an event when the preprocessor detects a reserved Modbus function code.

Note regarding the use of the Modbus preprocessor that if your network does not contain any Modbus-enabled devices, you should not enable this preprocessor in a network analysis policy that you apply to traffic.

You can use the following procedure to modify the ports the Modbus preprocessor monitors.

To configure the Modbus preprocessor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **Modbus Configuration** under **SCADA Preprocessors** is enabled:

- If the configuration is enabled, click **Edit**.

- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Modbus Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 Optionally, modify the **Ports** that the preprocessor inspects for Modbus traffic. You can specify an integer from 0 to 65535. Use commas to separate multiple ports.

Step 10 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Configuring the DNP3 Preprocessor

License: Protection

The Distributed Network Protocol (DNP3) is a SCADA protocol that was originally developed to provide consistent communication between electrical stations. DNP3 has also become widely used in the water, waste, transportation, and many other industries.

The DNP3 preprocessor detects anomalies in DNP3 traffic and decodes the DNP3 protocol for processing by the rules engine, which uses DNP3 keywords to access certain protocol fields. See DNP3 Keywords, page 30-76 for more information.

You must enable the DNP3 preprocessor rules in the following table if you want these rules to generate events. See Setting Rule States, page 27-19 for information on enabling rules.

Table 23-2 DNP3 Preprocessor Rules

Preprocessor Rule GID:SID	Description
145:1	When **Log bad CRC** is enabled, generates an event when the preprocessor detects a link layer frame with an invalid checksum.
145:2	Generates an event and blocks the packet when the preprocessor detects a DNP3 link layer frame with an invalid length.
145:3	Generates an event and blocks the packet during reassembly when the preprocessor detects a transport layer segment with an invalid sequence number.
145:4	Generates an event when the DNP3 reassembly buffer is cleared before a complete fragment can be reassembled. This happens when a segment carrying the FIR flag appears after other segments have been queued.
145:5	Generates an event when the preprocessor detects a DNP3 link layer frame that uses a reserved address.
145:6	Generates an event when the preprocessor detects a DNP3 request or response that uses a reserved function code.

Note regarding the use of the DNP3 preprocessor that, if your network does not contain any DNP3-enabled devices, you should not enable this preprocessor in a network analysis policy that you apply to traffic. See Configuring TCP Stream Preprocessing, page 24-28 for more information.

The following list describes the DNP3 preprocessor options you can configure.

Ports

Enables inspection of DNP3 traffic on each specified port. You can specify a single port or a comma-separated list of ports. You can specify a value from 0 to 65535 for each port.

Log bad CRCs

When enabled, validates the checksums contained in DNP3 link layer frames. Frames with invalid checksums are ignored.

You can enable rule 145:1 to generate events when invalid checksums are detected.

To configure the DNP3 preprocessor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **DNP3 Configuration** under **SCADA Preprocessors** is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The DNP3 Configuration page appears. A message at the bottom of the page identifies the network analysis policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 Optionally, modify the **Ports** that the preprocessor inspects for DNP3 traffic. You can specify an integer from 0 to 65535. Use commas to separate multiple ports.

Step 10 Optionally, select or clear the **Log bad CRCs** check box to specify whether to validate the checksums contained in DNP3 link layer frames and ignore frames with invalid checksums.

Step 11 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See the Network Analysis Policy Editing Actions table for more information.



Configuring Transport & Network Layer Preprocessing

You configure most transport at network layer preprocessors in a network analysis policy, which prepares traffic for inspection using the rules enabled in an intrusion policy. See Understanding Network Analysis and Intrusion Policies, page 18-1 for more information.

Transport and network layer preprocessors detect attacks that exploit IP fragmentation, checksum validation, and TCP and UDP session preprocessing. Before packets are sent to preprocessors, the packet decoder converts packet headers and payloads into a format that can be easily used by the preprocessors and the intrusion rules engine and detects various anomalous behaviors in packet headers. After packet decoding and before sending packets to other preprocessors, the inline normalization preprocessor normalizes traffic for inline deployments.

You can tailor transport and network layer preprocessor settings that you configure in network analysis policies by zone or network. Some transport and network layer settings apply globally to all traffic, and you configure these in an access control policy.

- Configuring Advanced Transport/Network Settings, page 24-1
- Verifying Checksums, page 24-4
- Normalizing Inline Traffic, page 24-6
- Defragmenting IP Packets, page 24-11
- Understanding Packet Decoding, page 24-16
- Using TCP Stream Preprocessing, page 24-20
- Using UDP Stream Preprocessing, page 24-31

Configuring Advanced Transport/Network Settings

License: Protection

Advanced transport and network preprocessor settings apply globally to all networks and zones where you apply your access control policy. You configure these advanced settings in an access control policy rather than in a network analysis policy.

The following sections describe these settings:

- Initiating Active Responses with Intrusion Drop Rules, page 24-2
- Troubleshooting: Logging Session Termination Messages, page 24-3

Initiating Active Responses with Intrusion Drop Rules

License: Protection

A drop rule is an intrusion rule or preprocessor rule whose rule state is set to Drop and Generate Events. In an inline deployment, the system responds to TCP or UDP drop rules by dropping the triggering packet and blocking the session where the packet originated. In a passive deployment, the system cannot drop the packet and does not block the session except with the use of active responses.

Tip Because UDP data streams are not typically thought of in terms of *sessions*, see Using UDP Stream Preprocessing, page 24-31 for further explanation of how the stream preprocessor uses the source and destination IP address fields in the encapsulating IP datagram header and the port fields in the UDP header to determine the direction of flow and identify a UDP session.

You can configure the **Maximum Active Responses** option to initiate one or more *active responses* to more precisely and specifically close a TCP connection or UDP session when an offending packet triggers a TCP or UDP drop rule.

When active responses are enabled in an inline deployment, the system responds to TCP drop rules by dropping the triggering packet and inserting a TCP Reset (RST) packet in both the client and server traffic. The system cannot drop the packet in a passive deployment; when active responses are enabled in a passive deployment, the system responds to TCP drop rules by sending a TCP reset to both the client and server ends of a TCP connection. When active responses are enabled in inline or passive deployments, the system closes a UDP session by sending an ICMP unreachable packet to each end of the session. Active responses are most effective in inline deployments because resets are more likely to arrive in time to affect the connection or session.

Depending on how you configure the **Maximum Active Responses** option, the system can also initiate additional active responses if it sees additional traffic from either end of the connection or session. The system initiates each additional active response, up to a specified maximum, after a specified number of seconds have elapsed since the previous response.

See Selecting The TCP Global Option, page 24-21 for information on setting the maximum number of active responses.

Note that a triggered **resp** or **react** rule also initiates an active response regardless of the configuration of **Maximum Active Responses**; however, **Maximum Active Responses** control whether the system initiates additional active responses for **resp** and **react** rules in the same way it controls the maximum number of active responses for drop rules. See Initiating Active Responses with Rule Keywords, page 30-84 for more information.

You can also use the `config response` command to configure the active response interface to use and the number of TCP resets to attempt in a passive deployment. See Setting the Active Response Reset Attempts and Interface, page 30-86 for more information.

No preprocessor rules are associated with the following options.

Maximum Active Responses

Specifies a maximum of 1 to 25 active responses per TCP connection. When additional traffic occurs on a connection where an active response has been initiated, and the traffic occurs more than **Minimum Response Seconds** after a previous active response, the system sends another active response unless the specified maximum has been reached. A setting of 0 disables active responses triggered by drop rules and disables additional active responses triggered by **resp** or **react** rules. For more information, see Initiating Active Responses with Intrusion Drop Rules, page 24-2 and Initiating Active Responses with Rule Keywords, page 30-84.

Minimum Response Seconds

Until **Maximum Active Responses** occur, specifies waiting 1 to 300 seconds before any additional traffic on a connection where the system has initiated an active response results in a subsequent active response.

To initiate active responses with drop rules:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 7 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 8 Click the edit icon ( ) next to **Transport/Network Layer Preprocessor Settings**.

The Transport/Network Layer Preprocessor Settings pop-up window appears.

Step 9 You have the following options:

- Specify a value 1 to 25 of **Maximum Active Responses** per TCP connection. A setting of 0 disables active responses triggered by drop rules and disables additional active responses triggered by **resp** or **react** rules.
- Specify a value 1 to 300 of **Minimum Response Seconds** to wait until either **Maximum Active Responses** occur or any additional traffic on a connection where the system has initiated an active response results in a subsequent active response.

Step 10 Click **OK.**

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Troubleshooting: Logging Session Termination Messages

License: Protection

Support might ask you during a troubleshooting call to configure your system to log a message when an individual connection exceeds the specified threshold. Changing the setting for this option will affect performance and should be done only with Support guidance.

To log session termination messages:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 7 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 8 Click the edit icon ( ) next to **Transport/Network Layer Preprocessor Settings**.

The Transport/Network Layer Preprocessor Settings pop-up window appears.

Step 9 Expand **Troubleshooting Options**.

Step 10 Specify for **Session Termination Logging Threshold** the number of bytes that result in a logged message when the session terminates and the specified number was exceeded.

The upper limit of 1GB is also restricted by the amount of memory on the device allocated for stream processing.

Step 11 Click **OK.**

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Verifying Checksums

License: Protection

The system can verify all protocol-level checksums to ensure that complete IP, TCP, UDP, and ICMP transmissions are received and that, at a basic level, packets have not been tampered with or accidentally altered in transit. A checksum uses an algorithm to verify the integrity of a protocol in the packet. The packet is considered to be unchanged if the system computes the same value that is written in the packet by the end host.

Disabling checksum verification may leave your network susceptible to insertion attacks. Note that the system does not generate checksum verification events. In an inline deployment, you can configure the system to drop packets with invalid checksums.

To configure checksum verifications:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Edit Policy page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **Checksum Verification** under Transport/Network Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Checksum Verification page appears. A message at the bottom of the page identifies the policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 You can set any of the options in the Checksum Verification section to **Enabled** or **Disabled** in a passive or inline deployment, or to **Drop** in an inline deployment:

- **ICMP Checksums**
- **IP Checksums**
- **TCP Checksums**
- **UDP Checksums**

Note that to drop offending packets, in addition to setting an option to **Drop** you must also enable **Inline Mode** in the associated network analysis policy. See Allowing Preprocessors to Affect Traffic in Inline Deployments, page 21-5 for more information. Note also that setting these options to **Drop** in a passive deployment is the same as setting them to **Enabled**.

Step 10 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Normalizing Inline Traffic

License: Protection

The inline normalization preprocessor normalizes traffic to minimize the chances of attackers evading detection in inline deployments. If you enable the inline normalization preprocessor in a network analysis policy, the system tests the following two conditions to ensure that you are using an inline deployment:

- **Inline Mode** is enabled in the policy. See Allowing Preprocessors to Affect Traffic in Inline Deployments, page 21-5.
- The access control policy where inline normalization is enabled is applied to a device that is deployed inline.

The preprocessor normalizes specified traffic only when both conditions are met.

You can specify normalization of any combination of IPv4, IPv6, ICMPv4, ICMPv6, and TCP traffic. Most normalizations are on a per-packet basis and are conducted by the inline normalization preprocessor. However, the TCP stream preprocessor handles most state-related packet and stream normalizations, including TCP payload normalization.

Inline normalization takes place immediately after decoding by the packet decoder and before processing by other preprocessors. Normalization proceeds from the inner to outer packet layers.

The inline normalization preprocessor does not generate events; it prepares packets for use by other preprocessors and the rules engine in inline deployments. The preprocessor also helps ensure that the packets the system processes are the same as the packets received by the hosts on your network.

Tip In an inline deployment, Cisco recommends that you configure the inline normalization preprocessor with the **Normalize TCP Payload** option enabled. In a passive deployment, Cisco recommends that you configure adaptive profiles. For more information, see Tuning Preprocessing in Passive Deployments, page 25-1.

Minimum TTL

When **Reset TTL** is greater than or equal to the value 1 to 255 set for this option, specifies the following:

- the minimum value the system will permit in the IPv4 Time to Live (TTL) field when **Normalize IPv4** is enabled; a lower value results in normalizing the packet value for TTL to the value set for **Reset TTL**
- the minimum value the system will permit in the IPv6 Hop Limit field when **Normalize IPv6** is enabled; a lower value results in normalizing the packet value for Hop Limit to the value set for **Reset TTL**

  The system assumes a value of 1 when the field is empty.

  Note that you can enable the following rules in the decoder rule category to generate events for this option:

- You can enable rule 116:428 to generate an event when the system detects an IPv4 packet with a TTL less than the specified minimum.
- You can enable rule 116:270 to generate an event when the system detects an IPv6 packet with a hop limit that is less than the specified minimum.

  See the packet decoder **Detect Protocol Header Anomalies** option in Configuring Packet Decoding, page 24-19 for more information.

Reset TTL

When set to a value 1 to 255 that is greater than or equal to **Minimum TTL**, normalizes the following:

- the IPv4 TTL field when **Normalize IPv4** is enabled
- the IPv6 Hop Limit field when **Normalize IPv6** is enabled

The system normalizes the packet by changing its TTL or Hop Limit value to the value set for this option when the packet value is less than **Minimum TTL**. Setting this option to a value of 0, or any value less than **Minimum TTL**, disables the option. The system assumes a value of 0 when the field is empty.

Normalize IPv4

Enables normalization of IPv4 traffic. The system also normalizes the TTL field as needed when this option is enabled and the value set for **Reset TTL** enables TTL normalization. You can also enable **Normalize Don't Fragment Bits** and **Normalize Reserved Bits** when this option is enabled.

When you enable this option, the system performs the following base IPv4 normalizations:

- truncates packets with excess payload to the datagram length specified in the IP header
- clears the Differentiated Services (DS) field, formerly known as the Type of Service (TOS) field
- sets all option octets to 1 (No Operation)

Normalize Don't Fragment Bit

Clears the single-bit Don't Fragment subfield of the IPv4 Flags header field. Enabling this option allows a downstream router to fragment packets if necessary instead of dropping them; enabling this option can also prevent evasions based on crafting packets to be dropped. You must enable **Normalize IPv4** to select this option.

Normalize Reserved Bit

Clears the single-bit Reserved subfield of the IPv4 Flags header field. You would typically enable this option. You must enable **Normalize IPv4** to select this option.

Normalize TOS Bit

Clears the one byte Differentiated Services field, formerly known as Type of Service. You must enable **Normalize IPv4** to select this option.

Normalize Excess Payload

Truncates packets with excess payload to the datagram length specified in the IP header plus the Layer 2 (for example, Ethernet) header, but does not truncate below the minimum frame length. You must enable **Normalize IPv4** to select this option.

Normalize IPv6

Sets all Option Type fields in the Hop-by-Hop Options and Destination Options extension headers to 00 (Skip and continue processing). The system also normalizes the Hop Limit field as needed when this option is enabled and the value set for **Reset TTL** enables hop limit normalization.

Normalize ICMPv4

Clears the 8-bit Code field in Echo (Request) and Echo Reply messages in ICMPv4 traffic.

Normalize ICMPv6

Clears the 8-bit Code field in Echo (Request) and Echo Reply messages in ICMPv6 traffic.

Normalize/Clear Reserved Bits

Clears the Reserved bits in the TCP header.

Normalize/Clear Option Padding Bytes

Clears any TCP option padding bytes.

Clear Urgent Pointer if URG=0

Clears the 16-bit TCP header Urgent Pointer field if the urgent (URG) control bit is not set.

Clear Urgent Pointer/URG on Empty Payload

Clears the TCP header Urgent Pointer field and the URG control bit if there is no payload.

Clear URG if Urgent Pointer is Not Set

Clears the TCP header URG control bit if the urgent pointer is not set.

Normalize Urgent Pointer

Sets the two-byte TCP header Urgent Pointer field to the payload length if the pointer is greater than the payload length.

Normalize TCP Payload

Enables normalization of the TCP Data field to ensure consistency in retransmitted data. Any segments that cannot be properly reassembled are dropped.

Remove Data on SYN

Removes data in synchronization (SYN) packets if your TCP operating system policy is **not** Mac OS.

This option also disables event generation for rule 129:2.

Remove Data on RST

Removes any data from a TCP reset (RST) packet.

Trim Data to Window

Trims the TCP Data field to the size specified in the Window field.

Trim Data to MSS

Trims the TCP Data field to the Maximum Segment Size (MSS) if the payload is longer than MSS.

Block Unrecoverable TCP Header Anomalies

When you enable this option, the system blocks anomalous TCP packets that, if normalized, would be invalid and likely would be blocked by the receiving host. For example, the system blocks any SYN packet transmitted subsequent to an established session.

The system also drops any packet that matches any of the following TCP stream preprocessor rules, regardless of whether the rules are enabled:

- 129:1
- 129:3
- 129:4
- 129:6

– 129:8
– 129:11
– 129:14 through 129:19

The Total Blocked Packets performance graph tracks the number of packets blocked in inline deployments and, in passive deployments, the number that would have been blocked in an inline deployment.

Explicit Congestion Notification

Enables per-packet or per-stream normalization of Explicit Congestion Notification (ECN) flags as follows:

- select **Packet** to clear ECN flags on a per-packet basis regardless of negotiation
- select **Stream** to clear ECN flags on a per-stream basis if ECN use was not negotiated

If you select **Stream**, you must also ensure that the TCP stream preprocessor **Require TCP 3-Way Handshake** option is enabled for this normalization to take place; see Selecting TCP Policy Options, page 24-22 for more information.

Allow These TCP Options

Disables normalization of specific TCP options you allow in traffic.

The system does not normalize options that you explicitly allow. It normalizes options that you do not explicitly allow by setting the options to No Operation (TCP Option 1).

The system always allows the Maximum Segment Size (MSS), Window Scale, and Time Stamp TCP options because these options are commonly used for optimal TCP performance. The system normalizes these commonly used options regardless of the configuration of **Allow These TCP Options**. The system does not automatically allow other less commonly used options.

You can allow specific options by configuring a comma-separated list of option keywords, option numbers, or both as shown in the following example:

```
sack, echo, 19
```

Specifying an option keyword is the same as specifying the number for one or more TCP options associated with the keyword. For example, specifying `sack` is the same as specifying TCP options 4 (Selective Acknowledgment Permitted) and 5 (Selective Acknowledgment). Option keywords are not case sensitive.

You can also specify `any`, which allows all TCP options and effectively disables normalization of all TCP options.

The following table summarizes how you can specify TCP options to allow. If you leave the field empty, the system allows only the MSS, Window Scale, and Time Stamp options.

Specify...	To allow...
sack	TCP options 4 (Selective Acknowledgment Permitted) and 5 (Selective Acknowledgment)
echo	TCP options 6 (Echo Request) and 7 (Echo Reply)
partial_order	TCP options 9 (Partial Order Connection Permitted) and 10 (Partial Order Service Profile)
conn_count	TCP Connection Count options 11 (CC), 12 (CC.New), and 13 (CC.Echo)

Specify...	To allow...
alt_checksum	TCP options 14 (Alternate Checksum Request) and 15 (Alternate Checksum)
md5	TCP option 19 (MD5 Signature)
the option number, 2 to 255	a specific option, including options for which there is no keyword
any	all TCP options; this setting effectively disables TCP option normalization

When you do not specify `any` for this option, normalizations include the following:

- except MSS, Window Scale, Time Stamp, and any explicitly allowed options, sets all option bytes to No Operation (TCP Option 1)
- sets the Time Stamp octets to No Operation if Time Stamp is present but invalid, or valid but not negotiated
- blocks the packet if Time Stamp is negotiated but not present
- clears the Time Stamp Echo Reply (TSecr) option field if the Acknowledgment (ACK) control bit is not set
- sets the MSS and Window Scale options to No Operation (TCP Option 1) if the SYN control bit is not set

To configure the inline normalizations preprocessor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Edit Policy page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices depending on whether **Inline Normalization** is enabled under Transport/Network Layer Preprocessors:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Inline Normalization page appears. A message at the bottom of the page identifies the policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 You can set any of the options described in Normalizing Inline Traffic, page 24-6.

Step 10 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Defragmenting IP Packets

License: Protection

When an IP datagram is broken into two or more smaller IP datagrams because it is larger than the maximum transmission unit (MTU), it is *fragmented*. A single IP datagram fragment may not contain enough information to identify a hidden attack. Attackers may attempt to evade detection by transmitting attack data in fragmented packets. The IP defragmentation preprocessor reassembles fragmented IP datagrams before the rules engine executes rules against them so the rules can more appropriately identify attacks in those packets. If fragmented datagrams cannot be reassembled, rules do not execute against them.

Note that you must enable IP defragmentation preprocessor rules, which have a generator ID (GID) of 123, if you want these rules to generate events. See Setting Rule States, page 27-19 for more information.

See the following sections for more information:

- Understanding IP Fragmentation Exploits, page 24-11
- Target-Based Defragmentation Policies, page 24-12
- Selecting Defragmentation Options, page 24-13
- Configuring IP Defragmentation, page 24-14

Understanding IP Fragmentation Exploits

License: Protection

Enabling IP defragmentation helps you detect attacks against hosts on your network, like the teardrop attack, and resource consumption attacks against the system itself, like the Jolt2 attack.

The Teardrop attack exploits a bug in certain operating systems that causes them to crash when trying to reassemble overlapping IP fragments. When enabled and configured to do so, the IP defragmentation preprocessor identifies the overlapping fragments. The IP defragmentation preprocessor detects the first packets in an overlapping fragment attack such as Teardrop, but does not detect subsequent packets for the same attack.

The Jolt2 attack sends a large number of copies of the same fragmented IP packet in an attempt to overuse IP defragmentors and cause a denial of service attack. A memory usage cap disrupts this and similar attacks in the IP defragmentation preprocessor, and places the system self-preservation above exhaustive inspection. The system is not overwhelmed by the attack, remains operational, and continues to inspect network traffic.

Different operating systems reassemble fragmented packets in different ways. Attackers who can determine which operating systems your hosts are running can also fragment malicious packets so that a target host reassembles them in a specific manner. Because the system does not know which operating systems the hosts on your monitored network are running, the preprocessor may reassemble and inspect the packets incorrectly, thus allowing an exploit to pass through undetected. To mitigate this kind of attack, you can configure the defragmentation preprocessor to use the appropriate method of defragmenting packets for each host on your network. See Target-Based Defragmentation Policies, page 24-12 for more information.

Note that you can also use adaptive profiles to dynamically select target-based policies for the IP defragmentation preprocessor using host operating system information for the target host in a packet. For more information, see Tuning Preprocessing in Passive Deployments, page 25-1.

Target-Based Defragmentation Policies

License: Protection

A host's operating system uses three criteria to determine which packet fragments to favor when reassembling the packet: the order in which the fragment was received by the operating system, its offset (the fragment's distance, in bytes, from the beginning of the packet), and its beginning and ending position compared to overlap fragments. Although every operating system uses these criteria, different operating systems favor different fragments when reassembling fragmented packets. Therefore, two hosts with different operating systems on your network could reassemble the same overlapping fragments in entirely different ways.

An attacker, aware of the operating system of one of your hosts, could attempt to evade detection and exploit that host by sending malicious content hidden in overlapping packet fragments. This packet, when reassembled and inspected, seems innocuous, but when reassembled by the target host, contains a malicious exploit. However, if you configure the IP defragmentation preprocessor to be aware of the operating systems running on your monitored network segment, it will reassemble the fragments the same way that the target host does, allowing it to identify the attack.

You can configure the IP defragmentation preprocessor to use one of seven defragmentation policies, depending on the operating system of the target host. The following table lists the seven policies and the operating systems that use each one. The First and Last policy names reflect whether those policies favor original or subsequent overlapping packets.

Table 24-1 Target-Based Defragmentation Policies

Policy	Operating Systems
BSD	AIX
	FreeBSD
	IRIX
	VAX/VMS
BSD-right	HP JetDirect

Table 24-1 Target-Based Defragmentation Policies (continued)

Policy	Operating Systems
First	Mac OS
	HP-UX
Linux	Linux
	OpenBSD
Last	Cisco IOS
Solaris	SunOS
Windows	Windows

Selecting Defragmentation Options

License: Protection

You can choose to simply enable or disable IP defragmentation; however, Cisco recommends that you specify the behavior of the enabled IP defragmentation preprocessor at a more granular level.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

You can configure the global **Preallocated Fragments** option:

Preallocated Fragments

The maximum number of individual fragments that the preprocessor can process at once. Specifying the number of fragment nodes to preallocate enables static memory allocation.

Caution Processing an individual fragment uses approximately 1550 bytes of memory. If the preprocessor requires more memory to process the individual fragments than the predetermined allowable memory limit for the device, the memory limit for the device takes precedence.

You can configure the following options for each IP defragmentation policy:

Networks

The IP address of the host or hosts to which you want to apply the defragmentation policy.

You can specify a single IP address or address block, or a comma-separated list of either or both. You can specify up to 255 total profiles, including the default policy. For information on using IPv4 and IPv6 address blocks in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or CIDR block/prefix length for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Note also that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks, and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

Policy

The defragmentation policy you want to use for a set of hosts on your monitored network segment. You can choose among seven policies: BSD, BSD-Right, First, Linux, Last, Solaris, and Windows. See Target-Based Defragmentation Policies, page 24-12 for detailed information on these policies.

Timeout

Specifies the maximum amount of time, in seconds, that the preprocessor engine can use when reassembling a fragmented packet. If the packet cannot be reassembled within the specified time period, the preprocessor engine stops attempting to reassemble the packet and discards received fragments.

Minimum TTL

Specifies the minimum acceptable TTL value a packet may have. This option detects TTL-based insertion attacks.

You can enable rule 123:1 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect Anomalies

Identifies fragmentation problems such as overlapping fragments.

You can enable the following rules to generate events for this option:

- 123:1 through 123:4
- 123:5 (BSD policy)
- 123:6 through 123:8

Overlap Limit

Specifies that when the configured number between 0 (unlimited) and 255 of overlapping segments in a session has been detected, defragmentation stops for that session. You must enable **Detect Anomalies** to configure this option. A blank value disables this option.

You can enable rule 123:12 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Minimum Fragment Size

Specifies that when a non-last fragment smaller than the configured number between 0 (unlimited) and 255 of bytes has been detected, the packet is considered malicious. You must enable **Detect Anomalies** to configure this option. A blank value disables this option.

You can enable rule 123:13 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Configuring IP Defragmentation

License: Protection

You can use the following procedure to configure the IP defragmentation preprocessor. For more information on the IP defragmentation preprocessor configuration options, see Selecting Defragmentation Options, page 24-13.

To configure IP defragmentation:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Edit Policy page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **IP Defragmentation** under Transport/Network Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The IP Defragmentation page appears. A message at the bottom of the page identifies the policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 Optionally, you can modify the setting for **Preallocated Fragments** in the Global Settings page area.

Step 10 You have two options:

- Add a new target-based policy. Click the add icon ( ) next to **Servers** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses in the **Host Address** field and click **OK**.

  You can specify a single IP address or address block, or a comma-separated list of either or both. You can create a total of 255 target-based policies including the default policy. For information on using IP address blocks in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

  Note that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks, and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

  A new entry appears in the list of targets on the left side of the page, highlighted to indicate that it is selected, and the Configuration section updates to reflect the current configuration for the policy you added.

- Modify the settings for an existing target-based policy. Click the configured address for a policy you have added under **Hosts** on the left side of the page, or click **default**.

  Your selection is highlighted and the Configuration section updates to display the current configuration for the policy you selected. To delete an existing target-based policy, click the delete icon ( ) next to the policy you want to remove.

Step 11 Optionally, you can modify any of the options in the Configuration page area.

Step 12 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Understanding Packet Decoding

License: Protection

Before sending captured packets to a preprocessor, the system first sends the packets to the packet decoder. The packet decoder converts packet headers and payloads into a format that preprocessors and the rules engine can easily use. Each stack layer is decoded in turn, beginning with the data link layer and continuing through the network and transport layers.

Note that you must enable packet decoder rules, which have a generator ID (GID) of 116, if you want these rules to generate events. See Setting Rule States, page 27-19 for more information.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Decode GTP Data Channel

Decodes the encapsulated GTP (General Packet Radio Service [GPRS] Tunneling Protocol) data channel. By default, the decoder decodes version 0 data on port 3386 and version 1 data on port 2152. You can use the `GTP_PORTS` default variable to modify the ports that identify encapsulated GTP traffic. See Optimizing Predefined Default Variables, page 2-13 for more information.

You can enable rules 116:297 and 116:298 to generate events for this option.

Detect Teredo on Non-Standard Ports

Inspects Teredo tunneling of IPv6 traffic that is identified on a UDP port other than port 3544.

The system always inspects IPv6 traffic when it is present. By default, IPv6 inspection includes the 4in6, 6in4, 6to4, and 6in6 tunneling schemes, and also includes Teredo tunneling when the UDP header specifies port 3544.

In an IPv4 network, IPv4 hosts can use the Teredo protocol to tunnel IPv6 traffic through an IPv4 Network Address Translation (NAT) device. Teredo encapsulates IPv6 packets within IPv4 UDP datagrams to permit IPv6 connectivity behind an IPv4 NAT device. The system normally uses UDP port 3544 to identify Teredo traffic. However, an attacker could use a non-standard port in an attempt to avoid detection. You can enable **Detect Teredo on Non-Standard Ports** to cause the system to inspect all UDP payloads for Teredo tunneling.

Teredo decoding occurs only on the first UDP header, and only when IPv4 is used for the outer network layer. When a second UDP layer is present after the Teredo IPv6 layer because of UDP data encapsulated in the IPv6 data, the rules engine uses UDP intrusion rules to analyze both the inner and outer UDP layers.

Note that intrusion rules 12065, 12066, 12067, and 12068 in the **policy-other** rule category detect, but do not decode, Teredo traffic. Optionally, you can use these rules to drop Teredo traffic in an inline deployment; however, you should ensure that these rules are disabled or set to generate events without dropping traffic when you enable **Detect Teredo on Non-Standard Ports**. See Filtering Rules in an Intrusion Policy, page 27-9 and Setting Rule States, page 27-19 for more information.

Detect Excessive Length Value

Detects when the packet header specifies a packet length that is greater than the actual packet length.

You can enable rules 116:6, 116:47, 116:97, and 116:275 to generate events for this option.

Detect Invalid IP Options

Detects invalid IP header options to identify exploits that use invalid IP options. For example, there is a denial of service attack against a firewall which causes the system to freeze. The firewall attempts to parse invalid Timestamp and Security IP options and fails to check for a zero length, which causes an irrecoverable infinite loop. The rules engine identifies the zero length option, and provides information you can use to mitigate the attack at the firewall.

You can enable rules 116:4 and 116:5 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect Experimental TCP Options

Detects TCP headers with experimental TCP options. The following table describes these options.

TCP Option	Description
9	Partial Order Connection Permitted
10	Partial Order Service Profile
14	Alternate Checksum Request
15	Alternate Checksum Data
18	Trailer Checksum
20	Space Communications Protocol Standards (SCPS)
21	Selective Negative Acknowledgements (SCPS)
22	Record Boundaries (SCPS)
23	Corruption (SPCS)
24	SNAP
26	TCP Compression Filter

Because these are experimental options, some systems do not account for them and may be open to exploits.

Note In addition to the experimental options listed in the above table, the system considers any TCP option with an option number greater than 26 to be experimental.

You can enable rule 116:58 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect Obsolete TCP Options

Detects TCP headers with obsolete TCP options. Because these are obsolete options, some systems do not account for them and may be open to exploits. The following table describes these options.

TCP Option	Description
6	Echo
7	Echo Reply
16	Skeeter
17	Bubba
19	MD5 Signature
25	Unassigned

You can enable rule 116:57 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect T/TCP

Detects TCP headers with the CC.ECHO option. The CC.ECHO option confirms that TCP for Transactions (T/TCP) is being used. Because T/TCP header options are not in widespread use, some systems do not account for them and may be open to exploits.

You can enable rule 116:56 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect Other TCP Options

Detects TCP headers with invalid TCP options not detected by other TCP decoding event options. For example, this option detects TCP options with the incorrect length or with a length that places the option data outside the TCP header.

You can enable rules 116:54, 116:55, and 116:59 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Detect Protocol Header Anomalies

Detects other decoding errors not detected by the more specific IP and TCP decoder options. For example, the decoder might detect a malformed data-link protocol header.

To generate events for this option, you can enable any packet decoder rule other than rules specifically associated with other packet decoder options. See Setting Rule States, page 27-19 for more information

Note that the following rules generate events triggered by anomalous IPv6 traffic: 116:270 through 116:274, 116:275 through 116:283, 116:291, 116:292, 116:295, 116:296, 116:406, 116:458, 116:460, 116:461.

Note also the following rules associated with the inline normalization preprocessor **Minimum TTL** option:

- You can enable rule 116:428 to generate an event when the system detects an IPv4 packet with a TTL less than the specified minimum.
- You can enable rule 116:270 to generate an event when the system detects an IPv6 packet with a hop limit that is less than the specified minimum.

See the inline normalization **Minimum TTL** option in Normalizing Inline Traffic, page 24-6 for more information.

Configuring Packet Decoding

License: Protection

You can configure packet decoding on the Packet Decoding configuration page. For more information on packet decoding configuration options, see Understanding Packet Decoding, page 24-16.

To configure packet decoding:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Edit Policy page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **Packet Decoding** under Transport/Network Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Packet Decoding page appears. A message at the bottom of the page identifies the policy layer that contains the configuration. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information

Step 9 You can enable or disable any of the detection options on the Packet Decoding page. See Understanding Packet Decoding, page 24-16 for more information.

Step 10 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Using TCP Stream Preprocessing

License: Protection

The TCP protocol defines various states in which connections can exist. Each TCP connection is identified by the source and destination IP addresses and source and destination ports. TCP permits only one connection with the same connection parameter values to exist at a time.

Note that you must enable TCP stream preprocessor rules, which have a generator ID (GID) of 129, if you want these rules to generate events. See Setting Rule States, page 27-19 for more information.

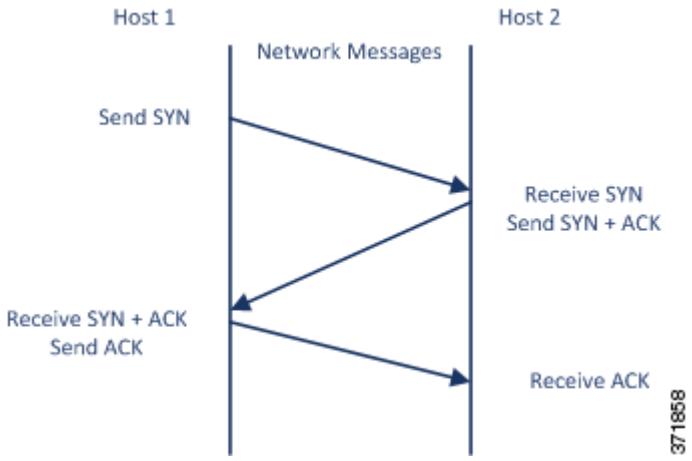
See the following sections for more information:

- Understanding State-Related TCP Exploits, page 24-20
- Initiating Active Responses with Intrusion Drop Rules, page 24-2
- Selecting The TCP Global Option, page 24-21
- Understanding Target-Based TCP Policies, page 24-21
- Selecting TCP Policy Options, page 24-22
- Reassembling TCP Streams, page 24-26
- Configuring TCP Stream Preprocessing, page 24-28

Understanding State-Related TCP Exploits

License: Protection

If you add the `flow` keyword with the `established` argument to an intrusion rule, the intrusion rules engine inspects packets matching the rule and the flow directive in stateful mode. Stateful mode evaluates only the traffic that is part of a TCP session established with a legitimate three-way handshake between a client and server. The following diagram illustrates a three-way handshake.



You can configure the system so that the preprocessor detects any TCP traffic that cannot be identified as part of an established TCP session, although this is not recommended for typical use because the events would quickly overload the system and not provide meaningful data.

Attacks like stick and snot use the system's extensive rule sets and packet inspection against itself. These tools generate packets based on the patterns in Snort-based intrusion rules, and send them across the network. If your rules do not include the `flow` or `flowbits` keyword to configure them for stateful inspection, each packet will trigger the rule, overwhelming the system. Stateful inspection allows you to ignore these packets because they are not part of an established TCP session and do not provide meaningful information. When performing stateful inspection, the rules engine detects only those attacks that are part of an established TCP session, allowing analysts to focus on these rather than the volume of events caused by stick or snot.

Selecting The TCP Global Option

License: Protection

The TCP stream preprocessor has one global option that controls how the TCP stream preprocessor functions.

No preprocessor rules are associated with this option.

Packet Type Performance Boost

Enables ignoring TCP traffic for all ports and application protocols that are not specified in enabled intrusion rules, except when a TCP rule with both the source and destination ports set to `any` has a `flow` or `flowbits` option. This performance improvement could result in missed attacks.

Understanding Target-Based TCP Policies

License: Protection

Different operating systems implement TCP in different ways. For example, Windows and some other operating systems require a TCP reset segment to have a precise TCP sequence number to reset a session, while Linux and other operating systems permit a range of sequence numbers. In this example, the stream preprocessor must understand exactly how the destination host will respond to the reset based on the sequence number. The stream preprocessor stops tracking the session only when the destination host considers the reset to be valid, so an attack cannot evade detection by sending packets after the preprocessor stops inspecting the stream. Other variations in TCP implementations include such things as whether an operating system employs a TCP timestamp option and, if so, how it handles the timestamp, and whether an operating system accepts or ignores data in a SYN packet.

Different operating systems also reassemble overlapping TCP segments in different ways. Overlapping TCP segments could reflect normal retransmissions of unacknowledged TCP traffic. They could also represent an attempt by an attacker, aware of the operating system of one of your hosts, to evade detection and exploit that host by sending malicious content hidden in overlapping segments. However, you can configure the stream preprocessor to be aware of the operating systems running on your monitored network segment so it reassembles segments the same way the target host does, allowing it to identify the attack.

You can create one or more TCP policies to tailor TCP stream inspection and reassembly to the different operating systems on your monitored network segment. For each policy, you identify one of thirteen operating system policies. You bind each TCP policy to a specific IP address or address block using as many TCP policies as you need to identify any or all of the hosts using a different operating system. The default TCP policy applies to any hosts on the monitored network that you do not identify in any other TCP policy, so there is no need to specify an IP address, CIDR block, or prefix length for the default TCP policy.

Note that you can also use adaptive profiles to dynamically select target-based policies for the TCP stream preprocessor using host operating system information for the target host in a packet. For more information, see Tuning Preprocessing in Passive Deployments, page 25-1.

The following table identifies the operating system policies and the host operating systems that use each.

Table 24-2 TCP Operating System Policies

Policy	Operating Systems
First	unknown OS
Last	Cisco IOS
BSD	AIX
	FreeBSD
	OpenBSD
Linux	Linux 2.4 kernel
	Linux 2.6 kernel
Old Linux	Linux 2.2 and earlier kernel
Windows	Windows 98
	Windows NT
	Windows 2000
	Windows XP
Windows 2003	Windows 2003
Windows Vista	Windows Vista
Solaris	Solaris OS
	SunOS
IRIX	SGI Irix
HPUX	HP-UX 11.0 and later
HPUX 10	HP-UX 10.2 and earlier
Mac OS	Mac OS 10 (Mac OS X)

Tip The First operating system policy could offer some protection when you do not know the host operating system. However, it may result in missed attacks. You should edit the policy to specify the correct operating system if you know it.

Selecting TCP Policy Options

License: Protection

The following list describes the options you can set to identify and control TCP traffic that the stream preprocessor inspects.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Network

Specifies the host IP addresses to which you want to apply the TCP stream reassembly policy.

You can specify a single IP address or address block. You can specify up to 255 total profiles including the default policy. For information on using IPv4 and IPv6 address blocks in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

Note that the `default` setting in the default policy specifies all IP addresses on your monitored network segment that are not covered by another target-based policy. Therefore, you cannot and do not need to specify an IP address or CIDR block/prefix length for the default policy, and you cannot leave this setting blank in another policy or use address notation to represent `any` (for example, 0.0.0.0/0 or ::/0).

Note also that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

Policy

Identifies the TCP policy operating system of the target host or hosts. If you select a policy other than **Mac OS**, the system removes the data from the synchronization (SYN) packets and disables event generation for rule 129:2.

For more information, see Understanding Target-Based TCP Policies, page 24-21.

Timeout

The number of seconds between 1 and 86400 the intrusion rules engine keeps an inactive stream in the state table. If the stream is not reassembled in the specified time, the intrusion rules engine deletes it from the state table.

Note If your device is deployed on a segment where the network traffic is likely to reach the device's bandwidth limits, you should consider setting this value higher (for example, to 600 seconds) to lower the amount of processing overhead.

Maximum TCP Window

Specifies the maximum TCP window size between 1 and 1073725440 bytes allowed as specified by a receiving host. Setting the value to 0 disables checking for the TCP window size.

Caution The upper limit is the maximum window size permitted by RFC, and is intended to prevent an attacker from evading detection, but setting a significantly large maximum window size could result in a self-imposed denial of service.

You can enable rule 129:6 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Overlap Limit

Specifies that when the configured number between 0 (unlimited) and 255 of overlapping segments in a session has been detected, segment reassembly stops for that session and, if **Stateful Inspection Anomalies** is enabled and the accompanying preprocessor rule is enabled, an event is generated.

You can enable rule 129:7 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Flush Factor

In an inline deployment, specifies that when a segment of decreased size has been detected subsequent to the configured number between 1 and 2048 of segments of non-decreasing size, the system flushes segment data accumulated for detection. Setting the value to 0 disables detection of this segment pattern, which can indicate the end of a request or response. Note that the Inline Normalization **Normalize TCP Payload** option must be enabled for this option the be effective. See Normalizing Inline Traffic, page 24-6 for more information.

Stateful Inspection Anomalies

Detects anomalous behavior in the TCP stack. When accompanying preprocessor rules are enabled, this may generate many events if TCP/IP stacks are poorly written.

You can enable the following rules to generate events for this option:

- 129:1 through 129:5
- 129:6 (Mac OS only)
- 129:8 through 129:11
- 129:13 through 129:19

See Setting Rule States, page 27-19 for more information:

TCP Session Hijacking

Detects TCP session hijacking by validating the hardware (MAC) addresses detected from both sides of a TCP connection during the 3-way handshake against subsequent packets received on the session. When the MAC address for one side or the other does not match, if **Stateful Inspection Anomalies** is enabled and one of the two corresponding preprocessor rules are enabled, the system generates events.

You can enable rules 129:9 and 129:10 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Consecutive Small Segments

When **Stateful Inspection Anomalies** is enabled, specifies a maximum number of 1 to 2048 consecutive small TCP segments allowed. Setting the value to 0 disables checking for consecutive small segments.

You must set this option together with the **Small Segment Size** option, either disabling both or setting a non-zero value for both. Note that receiving as many as 2000 consecutive segments, even if each segment was 1 byte in length, without an intervening ACK would be far more consecutive segments than you would normally expect.

You can enable rule 129:12 to generate events for this option. See Setting Rule States, page 27-19 for more information.

Small Segment Size

When **Stateful Inspection Anomalies** is enabled, specifies the 1 to 2048 byte TCP segment size that is considered small. Setting the value to 0 disables specifying the size of a small segment.

You must set this option together with the **Consecutive Small Segments** option, either disabling both or setting a non-zero value for both. Note that a 2048 byte TCP segment is larger than a normal 1500 byte Ethernet frame.

Ports Ignoring Small Segments

When **Stateful Inspection Anomalies**, **Consecutive Small Segments**, and **Small Segment Size** are enabled, optionally specifies a comma-separated list of one or more ports that ignore small TCP segment detection. Leaving this option blank specifies that no ports are ignored.

You can add any port to the list, but the list only affects ports specified in one of the **Perform Stream Reassembly on** port lists in the TCP policy.

Require TCP 3-Way Handshake

Specifies that sessions are treated as established only upon completion of a TCP three-way handshake. Disable this option to increase performance, protect from SYN flood attacks, and permit operation in a partially asynchronous environment. Enable it to avoid attacks that attempt to generate false positives by sending information that is not part of an established TCP session.

You can enable rule 129:20 to generate events for this option. See Setting Rule States, page 27-19 for more information.

3-Way Handshake Timeout

Specifies the number of seconds between 0 (unlimited) and 86400 (twenty-four hours) by which a handshake must be completed when **Require TCP 3-Way Handshake** is enabled. You must enable **Require TCP 3-Way Handshake** to modify the value for this option.

Packet Size Performance Boost

Sets the preprocessor to not queue large packets in the reassembly buffer. This performance improvement could result in missed attacks. Disable this option to protect against evasion attempts using small packets of one to twenty bytes. Enable it when you are assured of no such attacks because all traffic is comprised of very large packets.

Legacy Reassembly

Sets the stream preprocessor to emulate the deprecated Stream 4 preprocessor when reassembling packets, which lets you compare events reassembled by the stream preprocessor to events based on the same data stream reassembled by the Stream 4 preprocessor.

Asynchronous Network

Specifies whether the monitored network is an asynchronous network, that is, a network where the system sees only half the traffic. When this option is enabled, the system does not reassemble TCP streams to increase performance.

Perform Stream Reassembly on Client Ports, Server Ports, Both Ports

Specifies for client ports, server ports, or both, a comma-separated list of ports to identify the traffic for the stream preprocessor to reassemble. See Selecting Stream Reassembly Options, page 24-26.

Perform Stream Reassembly on Client Services, Server Services, Both Services

Specifies for client services, server services, or both, services to identify in the traffic for the stream preprocessor to reassemble. See Selecting Stream Reassembly Options, page 24-26.

Troubleshooting Options: Maximum Queued Bytes

Support might ask you during a troubleshooting call to specify the amount of data that can be queued on one side of a TCP connection. A value of 0 specifies an unlimited number of bytes.

Caution Changing the setting for this troubleshooting option will affect performance and should be done only with Support guidance.

Troubleshooting Options: Maximum Queued Segments

Support might ask you during a troubleshooting call to specify the maximum number of bytes of data segments that can be queued on one side of a TCP connection. A value of 0 specifies an unlimited number of data segment bytes.

Caution Changing the setting for this troubleshooting option will affect performance and should be done only with Support guidance.

Reassembling TCP Streams

License: Protection

The stream preprocessor collects and reassembles all the packets that are part of a TCP session's server-to-client communication stream, client-to-server communication stream, or both. This allows the rules engine to inspect the stream as a single, reassembled entity rather than inspecting only the individual packets that are part of a given stream.

See the following sections for more information:

- Understanding Stream-Based Attacks, page 24-26
- Selecting Stream Reassembly Options, page 24-26

Understanding Stream-Based Attacks

License: Protection

Stream reassembly allows the rules engine to identify stream-based attacks, which it may not detect when inspecting individual packets. You can specify which communication streams the rules engine reassembles based on your network needs. For example, when monitoring traffic on your web servers, you may only want to inspect client traffic because you are much less likely to receive malicious traffic from your own web server.

Selecting Stream Reassembly Options

License: Protection

In each TCP policy, you can specify a comma-separated list of ports to identify the traffic for the stream preprocessor to reassemble. If adaptive profiles are enabled, you can also list services that identify traffic to reassemble, either as an alternative to ports or in combination with ports. See Tuning Preprocessing in Passive Deployments, page 25-1 for information on enabling and using adaptive profiles.

You can specify ports, services, or both. You can specify separate lists of ports for any combination of client ports, server ports, and both. You can also specify separate lists of services for any combination of client services, server services, and both. For example, assume that you wanted to reassemble the following:

- SMTP (port 25) traffic from the client

- FTP server responses (port 21)
- telnet (port 23) traffic in both directions

You could configure the following:

- For client ports, specify `23, 25`
- For server ports, specify `21, 23`

Or, instead, you could configure the following:

- For client ports, specify `25`
- For server ports, specify `21`
- For both ports, specify `23`

Additionally, consider the following example which combines ports and services and would be valid when adaptive profiles are enabled:

- For client ports, specify `23`
- For client services, specify `smtp`
- For server ports, specify `21`
- For server services, specify `telnet`

Although you can also specify `all` as the argument to provide reassembly for all ports, Cisco does **not** recommend setting ports to `all` because it may increase the amount of traffic inspected by this preprocessor and slow performance unnecessarily.

TCP reassembly automatically and transparently includes ports that you add to other preprocessors. However, if you do explicitly add ports to TCP reassembly lists that you have added to other preprocessor configurations, these additional ports are handled normally. This includes port lists for the following preprocessors:

- FTP/Telnet (server-level FTP)
- DCE/RPC
- HTTP Inspect
- SMTP
- Session Initiation Protocol
- POP
- IMAP
- SSL

Negating a port (for example, !77) can improve performance by preventing the TCP stream preprocessor from processing traffic for that port.

Note that reassembling additional traffic types (client, server, both) increases resource demands.

If no preprocessor rule is mentioned in the following descriptions, the option is not associated with a preprocessor rule.

Perform Stream Reassembly on Client Ports

Enables stream reassembly based on ports for the client side of the connection. In other words, it reassembles streams destined for web servers, mail servers, or other IP addresses typically defined by the IP addresses specified in $HOME_NET. Use this option when you expect malicious traffic to originate from clients.

Perform Stream Reassembly on Client Services

Enables stream reassembly based on services for the client side of the connection. Use this option when you expect malicious traffic to originate from clients.

This feature requires Protection and Control licenses.

Perform Stream Reassembly on Server Ports

Enables stream reassembly based on ports for the server side of the connection only. In other words, it reassembles streams originating from web servers, mail servers, or other IP addresses typically defined by the IP addresses specified in $EXTERNAL_NET. Use this option when you want to watch for server side attacks. You can disable this option by not specifying ports.

Perform Stream Reassembly on Server Services

Enables stream reassembly based on services for the server side of the connection only. Use this option when you want to watch for server side attacks. You can disable this option by not specifying services.

This feature requires Protection and Control licenses.

Perform Stream Reassembly on Both Ports

Enables stream reassembly based on ports for both the client and server side of the connection. Use this option when you expect that malicious traffic for the same ports may travel in either direction between clients and servers. You can disable this option by not specifying ports.

Perform Stream Reassembly on Both Services

Enables stream reassembly based on services for both the client and server side of the connection. Use this option when you expect that malicious traffic for the same services may travel in either direction between clients and servers.You can disable this option by not specifying services.

This feature requires Protection and Control licenses.

Configuring TCP Stream Preprocessing

License: Protection

You can configure TCP stream preprocessing, including TCP policies. For more information on the TCP stream preprocessor configuration options, see Selecting TCP Policy Options, page 24-22.

To configure the stream preprocessor to track TCP sessions:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Edit Policy page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **TCP Stream Configuration** under Transport/Network Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The TCP Stream Configuration page appears. A message at the bottom of the page identifies the policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 Optionally, modify **Packet Type Performance Boost** under Global Settings. See Selecting The TCP Global Option, page 24-21 for more information.

Step 10 You have two options:

- Add a new target-based policy. Click the add icon ( ) next to **Hosts** on the left side of the page. The Add Target pop-up window appears. Specify one or more IP addresses in the **Host Address** field and click **OK**.

  You can specify a single IP address or address block. You can create a total of 255 target-based policies including the default policy. For information on using IP address blocks in the ASA FirePOWER module, see IP Address Conventions, page 1-4.

  Note that for a target-based policy to process traffic, the networks you identify must match or be a subset of the networks and zones handled by the network analysis policy where you configure the target-based policy. See Customizing Preprocessing with Network Analysis Policies, page 20-2 for more information.

  A new entry appears in the list of targets on the left side of the page, highlighted to indicate that it is selected, and the Configuration section updates to reflect the current configuration for the policy you added.

- Modify the settings for an existing target-based policy. Click the configured address for a policy you have added under **Hosts** on the left side of the page, or click **default**.

  Your selection is highlighted and the Configuration section updates to display the current configuration for the policy you selected. To delete an existing target-based policy, click the delete icon ( ) next to the policy you want to remove.

Step 11 Optionally, modify any of the TCP policy options under Configuration.

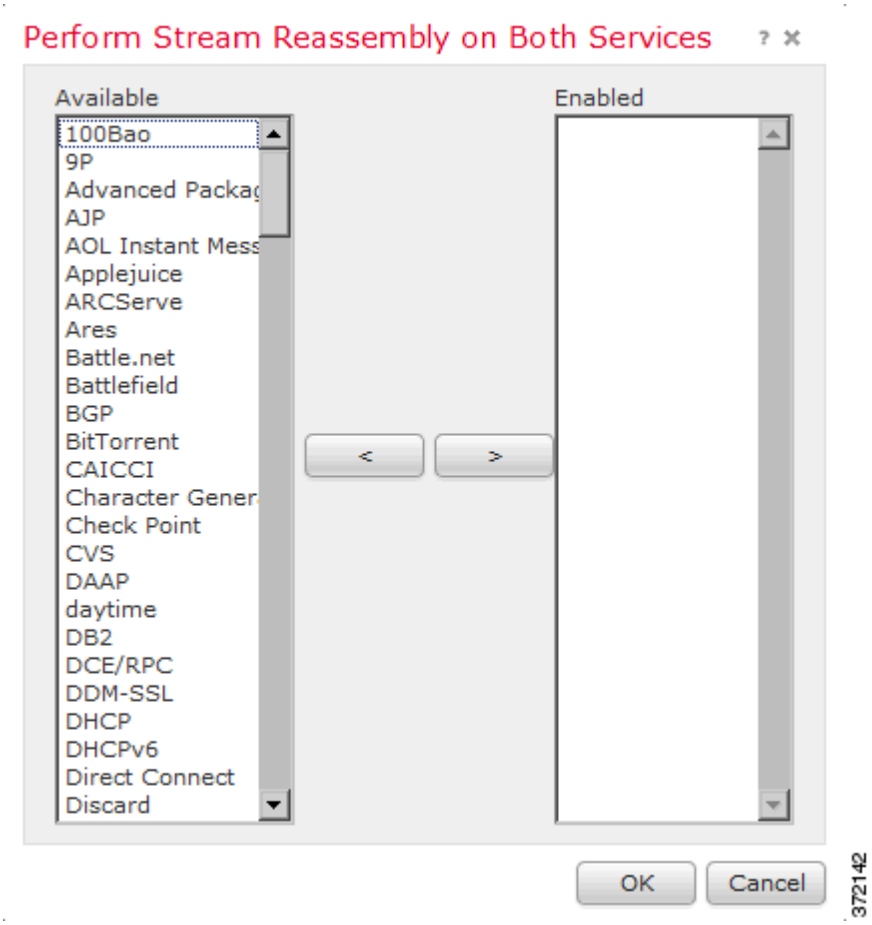
For specific instructions on modifying settings for stream reassembly based on client services, server services, or both go to step 12; otherwise, go to step 15.

For more information, see Selecting TCP Policy Options, page 24-22, and Selecting Stream Reassembly Options, page 24-26.

Step 12 To modify settings for stream reassembly based on client, server, or both services, click inside the field you want to modify or click **Edit** next to the field.

The pop-up window for the field you selected appears.

For example, the following graphic shows the Perform Stream Reassembly on Both Services pop-up window.



Note that you can enable adaptive profiles to monitor traffic for the stream preprocessor to reassemble based on services discovered on your network. See Tuning Preprocessing in Passive Deployments, page 25-1 for more information.

Step 13 You have two choices:

- To add services to monitor, select one or more services from the **Available** list on the left, then click the right arrow (>) button.
- To remove a service, select it from the **Enabled** list on the right, then click the left arrow (<) button.

Use Ctrl or Shift while clicking to select multiple service detectors. You can also click and drag to select multiple adjacent service detectors.

Step 14 Click **OK** to add the selections.

The TCP Stream Configuration page is displayed and the services are updated.

Step 15 Optionally, expand the **Troubleshooting Options** and modify either of the TCP stream preprocessing policy settings only if asked to do so by Support. For more information, see Selecting TCP Policy Options, page 24-22.

Step 16 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Using UDP Stream Preprocessing

License: Protection

UDP stream preprocessing occurs when the rules engine processes packets against a UDP rule that includes the `flow` keyword (see Applying Rules to a TCP or UDP Client or Server Flow, page 30-51) using any of the following arguments:

- `Established`
- `To Client`
- `From Client`
- `To Server`
- `From Server`

UDP is a connectionless protocol that does not provide a means for two endpoints to establish a communication channel, exchange data, and close the channel. UDP data streams are not typically thought of in terms of *sessions*. However, the stream preprocessor uses the source and destination IP address fields in the encapsulating IP datagram header and the port fields in the UDP header to determine the direction of flow and identify a session. A session ends when a configurable timer is exceeded, or when either endpoint receives an ICMP message that the other endpoint is unreachable or the requested service is unavailable.

Note that the system does not generate events related to UDP stream preprocessing; however, you can enable related packet decoder rules to detect UDP protocol header anomalies. For information on events generated by the packet decoder, see Understanding Packet Decoding, page 24-16.

Configuring UDP Stream Preprocessing

License: Protection

You can configure UDP stream preprocessing.

To configure the stream preprocessor to track UDP sessions:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Edit Policy page appears.

Step 7 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 8 You have two choices, depending on whether **UDP Stream Configuration** under Transport/Network Layer Preprocessors is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The UDP Stream Configuration page appears. A message at the bottom of the page identifies the policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 Optionally, configure a **Timeout** value to specify the number of seconds between 1 and 86400 the preprocessor keeps an inactive stream in the state table. If additional datagrams are not seen in the specified time, the preprocessor deletes the stream from the state table.

Step 10 Optionally, select **Packet Type Performance Boost** to ignore UDP traffic for all ports and application protocols that are not specified in enabled rules, except when a UDP rule with both the source and destination ports set to `any` has a `flow` or `flowbits` option. This performance improvement could result in missed attacks.

Step 11 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.



Tuning Preprocessing in Passive Deployments

Typically, the system uses the static settings in your network analysis policy to preprocess and analyze traffic. With the adaptive profiles feature, however, the system can adapt to network traffic by associating traffic with host information and processing the traffic accordingly.

When a host receives traffic, the operating system running on the host reassembles IP fragments. The order used for that reassembly depends on the operating system. Similarly, each operating system may implement TCP in different ways, and therefore reassemble TCP streams differently. If preprocessors reassemble data using a format other than that used for the operating system of the destination host, the system may miss content that could be malicious when reassembled on the receiving host.

Tip In a passive deployment, Cisco recommends that you configure adaptive profiles. In an inline deployment, Cisco recommends that you configure the inline normalization preprocessor with the **Normalize TCP Payload** option enabled. For more information, see Normalizing Inline Traffic, page 24-6.

For more information on using adaptive profiles to improve reassembly of packet fragments and TCP streams, see the following topics:

- Understanding Adaptive Profiles, page 25-1
- Configuring Adaptive Profiles, page 25-2

Understanding Adaptive Profiles

License: Protection

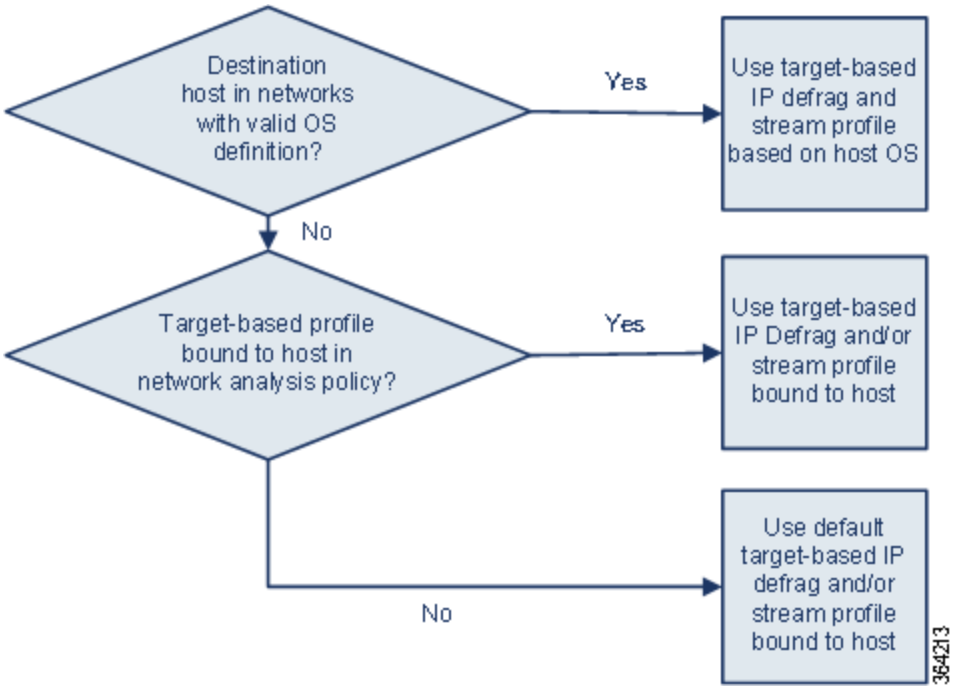
Adaptive profiles enable use of the most appropriate operating system profiles for IP defragmentation and TCP stream preprocessing. For more information on the aspects of the network analysis policy affected by adaptive profiles, see Defragmenting IP Packets, page 24-11 and Using TCP Stream Preprocessing, page 24-20.

Using Adaptive Profiles with Preprocessors

License: Protection

Adaptive profiles help to defragment IP packets and reassemble streams in the same way as the operating system on the target host. The intrusion rules engine then analyzes the data in the same format as that used by the destination host.

Adaptive profiles switch to the appropriate operating system profile based on the operating system in the host profile for the target host, as illustrated in the following diagram.



For example, you configure adaptive profiles for the 10.6.0.0/16 subnet and set the default IP Defragmentation target-based policy to Linux. The ASA FirePOWER module where you configure the settings includes the 10.6.0.0/16 subnet.

When a device detects traffic from Host A, which is not in the 10.6.0.0/16 subnet, it uses the Linux target-based policy to reassemble IP fragments. However, when it detects traffic from Host B, which is in the 10.6.0.0/16 subnet, it retrieves Host B's operating system data, where Host B is running Microsoft Windows XP Professional. The system uses the Windows target-based profile to do the IP defragmentation for the traffic destined for Host B.

See Defragmenting IP Packets, page 24-11 for information on the IP Defragmentation preprocessor. See Using TCP Stream Preprocessing, page 24-20 for information on the stream preprocessor.

Configuring Adaptive Profiles

License: Protection

To use host information to determine which target-based profiles are used for IP defragmentation and TCP stream preprocessing, you can configure adaptive profiles.

When you configure adaptive profiles, you need to bind the adaptive profile setting to a specific network or networks. To successfully use adaptive profiles, that network must be in the segment monitored by the device.

You can indicate the hosts in the network where adaptive profiles should be used to process traffic by specifying an IP address, a block of addresses, or a network variable with the desired value configured in the variable set linked to the default intrusion policy for your access control policy. See Setting the Default Intrusion Policy for Access Control, page 20-1 for more information.

You can use any of these addressing methods alone or in any combination as a list of IP addresses, address blocks, or variables separated by commas, as shown in the following example:

```
192.168.1.101, 192.168.4.0/24, $HOME_NET
```

For information on specifying address blocks, see IP Address Conventions, page 1-4.

Tip You can apply adaptive profiles to all hosts in the network by using a variable with a value of `any` or by specifying `0.0.0.0/0` as the network value.

To configure adaptive profiles:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy.**

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Detection Enhancement Settings**.

The Detection Enhancement Settings pop-up window appears.

Step 5 Select **Adaptive Profiles - Enabled** to enable adaptive profiles.

Step 6 Optionally, in the **Adaptive Profiles - Attribute Update Interval** field, type the number of minutes that should elapse between synchronization of data.

Note Increasing the value for this option could improve performance in a large network.

Step 7 In the **Adaptive Profiles - Networks** field, type the specific IP address, address block, or variable, or a list that includes any of these addressing methods separated by commas, to identify any host in the network for which you want to use adaptive profiles.

See Working with Variable Sets, page 2-13 for information on configuring variables.

Step 8 Click **OK** to retain your settings.



Getting Started with Intrusion Policies

Intrusion policies are defined sets of intrusion detection and prevention configurations that inspect traffic for security violations and, in inline deployments, can block or alter malicious traffic. Intrusion policies are invoked by your access control policy and are the system's last line of defense before traffic is allowed to its destination.

Cisco delivers several intrusion policies with the ASA FirePOWER module. By using system-provided policies you can take advantage of the experience of the Cisco Vulnerability Research Team (VRT). For these policies, the VRT sets intrusion and preprocessor rule states (enabled or disabled), as well as provides the initial configurations for other advanced settings. An enabled rule causes the system to generate intrusion events for (and optionally block) traffic matching the rule. Disabling a rule stops processing of the rule.

Tip System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules. Understanding Network Analysis and Intrusion Policies, page 18-1 provides an overview of how network analysis and intrusion policies work together to examine your traffic, as well as some basics on using the navigation panel, resolving conflicts, and committing changes.

If you create a custom intrusion policy, you can:

- Tune detection by enabling and disabling rules, as well as by writing and adding your own rules.
- Configure various advanced settings such as external alerting, sensitive data preprocessing, and global rule thresholding.
- Use layers as building blocks to efficiently manage multiple intrusion policies.

When tailoring your intrusion policy, especially when enabling and adding rules, keep in mind that some intrusion rules require that traffic first be decoded or preprocessed in a certain way. Before an intrusion policy examines a packet, the packet is preprocessed according to configurations in a network analysis policy. If you disable a required preprocessor, the system automatically uses it with its current settings, although the preprocessor remains disabled in the network analysis policy user interface.

Note Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task. For more information, see Limitations of Custom Policies, page 18-11.

After you configure a custom intrusion policy, you can use it as part of your access control configuration by associating the intrusion policy with one or more access control rules or an access control policy's default action. This forces the system to use the intrusion policy to examine certain allowed traffic before the traffic passes to its final destination. A variable set that you pair with the intrusion policy allows you to accurately reflect your home and external networks and, as appropriate, the servers on your network. For more information, see Controlling Traffic Using Intrusion and File Policies, page 11-1.

This chapter explains how to create a simple custom intrusion policy. The chapter also contains basic information on managing intrusion policies: editing, comparing, and so on. For more information, see:

- Creating a Custom Intrusion Policy, page 26-2
- Managing Intrusion Policies, page 26-3
- Editing Intrusion Policies, page 26-4
- Applying an Intrusion Policy, page 26-7
- Generating a Report of Current Intrusion Settings, page 26-8
- Comparing Two Intrusion Policies or Revisions, page 26-9

Creating a Custom Intrusion Policy

License: Protection

When you create a new intrusion policy you must give it a unique name, specify a base policy, and specify drop behavior.

The base policy defines the intrusion policy's default settings. Modifying a setting in the new policy overrides—but does not change—the settings in the base policy. You can use either a system-provided or custom policy as your base policy. For more information, see Understanding the Base Layer, page 19-2.

The intrusion policy's drop behavior, or **Drop when Inline** setting, determines how the system handles drop rules (intrusion or preprocessor rules whose rule state is set to Drop and Generate Events) and other intrusion policy configurations that affect traffic. You should enable drop behavior in inline deployments when you want to drop or replace malicious packets. Note that in passive deployments, the system cannot affect traffic flow regardless of the drop behavior. For more information, see Setting Drop Behavior in an Inline Deployment, page 26-5.

To create an intrusion policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy**.

The Intrusion Policy page appears.



Tip You can also import a policy from another ASA FirePOWER module; see Importing and Exporting Configurations, page B-1.

Step 2 Click **Create Policy**.

If you have unsaved changes in another policy, click **Cancel** when prompted to return to the Intrusion Policy page. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Create Intrusion Policy pop-up window appears.

Step 3 Give the policy a unique **Name** and, optionally, a **Description**.

Step 4 Specify the initial **Base Policy**.

You can use either a system-provided or custom policy as your base policy.

Caution Do **not** use `Experimental Policy 1` unless instructed to do so by a Cisco representative. Cisco uses this policy for testing.

Step 5 Set the system's drop behavior in an inline deployment:

- To allow intrusion policies to affect traffic and generate events, enable **Drop when Inline**.
- To prevent intrusion policies from affecting traffic while still generating events, disable **Drop when Inline**.

Step 6 Create the policy:

- Click **Create Policy** to create the new policy and return to the Intrusion Policy page. The new policy has the same settings as its base policy.
- Click **Create and Edit Policy** to create the policy and open it for editing in the advanced intrusion policy editor; see Editing Intrusion Policies, page 26-4.

Managing Intrusion Policies

License: Protection

On the Intrusion Policy page (**Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy**) you can view your current custom intrusion policies, along with the following information:

- the time and date the policy was last modified (in local time)
- whether the **Drop when Inline** setting is enabled, which allows you to drop and modify traffic in an inline deployment
- which access control policies are using the intrusion policy to inspect traffic
- whether a policy has unsaved changes

Options on the Intrusion Policy page allow you to take the actions in the following table.

Table 26-1 Intrusion Policy Management Actions

To...	You can...	See...
create a new intrusion policy	click **Create Policy**.	Creating a Custom Intrusion Policy, page 26-2
edit an existing intrusion policy	click the edit icon ( ).	Editing Intrusion Policies, page 26-4
reapply an intrusion policy	click the apply icon ( ).	Applying an Intrusion Policy, page 26-7
export an intrusion policy to import on another ASA FirePOWER module	click the export icon ( ).	Exporting Configurations, page B-1

Table 26-1 Intrusion Policy Management Actions (continued)

To...	You can...	See...
view a PDF report that lists the current configuration settings in a intrusion policy	click the report icon ( ).	Generating a Report of Current Intrusion Settings, page 26-8
compare the settings of two intrusion policies or two revisions of the same policy	click **Compare Policies**.	Comparing Two Intrusion Policies or Revisions, page 26-9
delete an intrusion policy	click the delete icon ( ), then confirm that you want to delete the policy. You cannot delete an intrusion policy if an access control policy references it.	

Editing Intrusion Policies

License: Protection

When you create a new intrusion policy, it has the same intrusion rule and advanced settings as its base policy. The following table explains the most common actions taken when editing an intrusion policy:

Table 26-2 Intrusion Policy Editing Actions

To...	You can...	See...
specify drop behavior in an inline deployment	select or clear the **Drop when Inline** check box on the Policy Information page.	Setting Drop Behavior in an Inline Deployment, page 26-5
change the base policy	select a base policy from the **Base Policy** drop-down list on the Policy Information page.	Changing the Base Policy, page 19-3
view the settings in the base policy	click **Manage Base Policy** on the Policy Information page	Understanding the Base Layer, page 19-2
display or configure intrusion rules	click **Manage Rules** on the Policy Information page.	Viewing Rules in an Intrusion Policy, page 27-2
display a filtered view of intrusion rules by current rule state and, optionally, configure those rules	on the Policy Information page, click **View** next to the number of rules under **Manage Rules** that are set to Generate Events or to Drop and Generate Events.	Filtering Rules in an Intrusion Policy, page 27-9
enable, disable, or edit advanced settings	click **Advanced Settings** in the navigation panel	Configuring Advanced Settings in an Intrusion Policy, page 26-6
manage policy layers	click **Policy Layers** in the navigation panel	Using Layers in a Network Analysis or Intrusion Policy, page 19-1

When tailoring an intrusion policy, especially when enabling and adding rules, keep in mind that some intrusion rules require that traffic first be decoded or preprocessed in a certain way. Before an intrusion policy examines a packet, the packet is preprocessed according to configurations in a network analysis policy. If you disable a required preprocessor, the system automatically uses it with its current settings, although the preprocessor remains disabled in the network analysis policy user interface.

Note Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task. For more information, see Limitations of Custom Policies, page 18-11.

The system caches one intrusion policy. While editing an intrusion policy, if you select any menu or other path to another page, your changes stay in the system cache even if you leave the page. In addition to the actions you can perform in the table above, Understanding Network Analysis and Intrusion Policies, page 18-1 provides information on resolving conflicts and committing changes

To edit a intrusion policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy**.

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the intrusion policy you want to configure.

The intrusion policy editor appears, focused on the Policy Information page and with a navigation panel on the left.

Step 3 Edit your policy. Take any of the actions summarized above.

Step 4 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Setting Drop Behavior in an Inline Deployment

License: Protection

In an inline deployment, an intrusion policy can block and modify traffic:

- *Drop rules* can drop matching packets and generate intrusion events. To configure an intrusion or preprocessor drop rule, set its state to Drop and Generate Events; see Setting Rule States, page 27-19.
- Intrusion rules can use the `replace` keyword to replace malicious content; see Replacing Content in Inline Deployments, page 30-29.

For intrusion rules to affect traffic, you must correctly configure drop rules and rules that replace content, as well as correctly deploy the system inline. Finally, you must enable the intrusion policy's *drop behavior*, or **Drop when Inline** setting.

Note To block the transfer of malware files over FTP, you must not only correctly configure network-based advanced malware protection (AMP), but also enable **Drop when Inline** in your access control policy's default intrusion policy. To determine or change the default intrusion policy, see Setting the Default

Intrusion Policy for Access Control, page 20-1.

If you want to assess how your configuration would function in an inline deployment without actually affecting traffic, you can disable drop behavior. In this case, the system generates intrusion events but does not drop packets that trigger drop rules. When you are satisfied with the results, you can enable drop behavior.

Note that in passive deployments the system cannot affect traffic regardless of the drop behavior. In other words, in a passive deployment, rules set to Drop and Generate Events behave identically to rules set to Generate Events—the system generates intrusion events but cannot drop packets.

When you view intrusion events, workflows can include the *inline result*, which indicates whether traffic was actually dropped, or whether it only would have dropped. When a packet matches a drop rule, the inline result is:

- `Dropped`, for packets dropped by a correctly configured inline deployment with drop behavior enabled
- `Would have dropped`, for packets that were not dropped either because your device is deployed passively or because drop behavior is disabled. Note that the inline result is always `Would have dropped` for packets seen while the system is pruning, regardless of deployment.

To set the drop behavior of an intrusion policy in an inline deployment:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy**.

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

The Policy Information page appears.

Step 3 Set the policy's drop behavior:

- To allow intrusion rules to affect traffic and generate events, enable **Drop when Inline**.
- To prevent intrusion rules from affecting traffic while still generating events, disable **Drop when Inline**.

Step 4 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Configuring Advanced Settings in an Intrusion Policy

License: Protection

An intrusion policy's *advanced settings* require specific expertise to configure. The base policy for your intrusion policy determines which advanced settings are enabled by default and the default configuration for each.

When you select **Advanced Settings** in the navigation panel of an intrusion policy, the policy lists its advanced settings by type. On the Advanced Settings page, you can enable or disable advanced settings in your intrusion policy, as well as access advanced setting configuration pages.

An advanced setting must be enabled for you to configure it. When you enable an advanced setting, a sublink to the configuration page for the advanced setting appears beneath the **Advanced Settings** link in the navigation panel, and an **Edit** link to the configuration page appears next to the advanced setting on the Advanced Settings page.

Tip To revert an advanced setting's configuration to the settings in the base policy, click **Revert to Defaults** on the configuration page for the advanced setting. When prompted, confirm that you want to revert.

When you disable an advanced setting, the sublink and **Edit** link no longer appear, but your configurations are retained. Note that some intrusion policy configurations (sensitive data rules, SNMP alerts for intrusion rules) require enabled and correctly configured advanced settings. You cannot save an intrusion policy misconfigured in this way; see Resolving Conflicts and Committing Policy Changes, page 18-15.

Modifying the configuration of an advanced setting requires an understanding of the configuration you are modifying and its potential impact on your network. The following sections provide links to specific configuration details for each advanced setting.

Specific Threat Detection

The sensitive data preprocessor detects sensitive data such as credit card numbers and Social Security numbers in ASCII text. For information on configuring this preprocessor, see Detecting Sensitive Data, page 28-19.

Note that other preprocessors that detect specific threats (back orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic) are configured in network analysis policies. For more information, see Detecting Specific Threats, page 28-1.

Intrusion Rule Thresholds

Global rule thresholding can prevent your system from being overwhelmed with a large number of events by allowing you to use thresholds to limit the number of times the system logs and displays intrusion events. For more information, see Globally Limiting Intrusion Event Logging, page 29-1.

External Responses

In addition to the various views of intrusion events within the user interface, you can enable logging to system log (syslog) facilities or send event data to an SNMP trap server. Per policy, you can specify intrusion event notification limits, set up intrusion event notification to external logging facilities, and configure external responses to intrusion events. For more information, see:

- Configuring SNMP Responses, page 39-3
- Configuring Syslog Responses, page 39-6

Applying an Intrusion Policy

License: Protection

After you apply an intrusion policy using access control (see Deploying Configuration Changes, page 4-11), you can reapply the intrusion policy at any time. This allows you to implement intrusion policy changes on your monitored network without reapplying the access control policy. While reapplying, you can also view a comparison report to review the changes made since the last time the intrusion policy was applied.

Note the following when reapplying intrusion policies:

- You can schedule intrusion policy reapply tasks to recur on a regular basis; see Automating Applying an Intrusion Policy, page 42-3.
- When you import a rule update, you can automatically apply intrusion policies after the import completes. If you do not enable this option, you must manually reapply the policies changed by the rule update. See Importing Rule Updates and Local Rule Files, page 46-9 for more information.

To reapply an intrusion policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy**.

The Intrusion Policy page appears.

Step 2 Click the apply icon ( ) next to the policy you want to reapply.

The Reapply Intrusion Policy window appears.

Step 3 Click **Reapply**.

The policy is reapplied. You can monitor the status of the apply using the task queue (**Monitoring > ASA FirePOWER Monitoring > Task Status**). See Viewing the Task Queue, page C-1 for more information.

Generating a Report of Current Intrusion Settings

License: Protection

An intrusion policy report is a record of the policy configuration at a specific point in time. The system combines the settings in the base policy with the settings of the policy layers, and makes no distinction between which settings originated in the base policy or policy layer.

You can use the report, which contains the following information, for auditing purposes or to inspect the current configuration.

Table 26-3 Intrusion Policy Report Sections

Section	Description
Policy Information	Provides the name and description of the intrusion policy, the name of the user who last modified the policy, and the date and time the policy was last modified. Also indicates whether dropping packets in an inline deployment is enabled or disabled, the current rule update version, and whether the base policy is locked to the current rule update.
Advanced Settings	Lists all enabled intrusion policy advanced settings and their configurations.
Rules	Provides a list of all enabled rules and their actions.

You can also generate a comparison report that compares two intrusion policies, or two revisions of the same policy. For more information, see Comparing Two Intrusion Policies or Revisions, page 26-9.

To view an intrusion policy report:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy**.

The Intrusion Policy page appears.

Step 2 Click the report icon ( ) next to the intrusion policy for which you want to generate a report. Remember to commit any potential changes before you generate an intrusion policy report; only committed changes appear in the report.

The system generates the intrusion policy report. You are prompted to save the report to your computer.

Comparing Two Intrusion Policies or Revisions

License: Protection

To review policy changes for compliance with your organization's standards or to optimize system performance, you can examine the differences between two intrusion policies. You can compare any two intrusion policies or two revisions of the same intrusion policy, for the intrusion policies you can access. Optionally, after you compare, you can then generate a PDF report to record the differences between the two policies or policy revisions.

There are two tools you can use to compare intrusion policies:

- The comparison view displays only the differences between two intrusion policies or intrusion policy revisions in a side-by-side format; the name of each policy appears in the title bar on the left and right sides of the comparison view.

  You can use this to view and navigate both policy revisions on the user interface, with their differences highlighted.

- The comparison report creates a record of only the differences between two intrusion policies or intrusion policy revisions in a format similar to the intrusion policy report, but in PDF format.

  You can use this to save, copy, print and share your policy comparisons for further examination.

For more information on understanding and using the intrusion policy comparison tools, see:

- Using the Intrusion Policy Comparison View, page 26-9
- Using the Intrusion Policy Comparison Report, page 26-10

Using the Intrusion Policy Comparison View

License: Protection

The comparison view displays both intrusion policies or policy revisions in a side-by-side format, with each policy or policy revision identified by name in the title bar on the left and right sides of the comparison view. The time of last modification and the last user to modify are displayed to the right of the policy name. Note that the Intrusion Policy page displays the time a policy was last modified in local time, but the intrusion policy report lists the time modified in UTC. Differences between the two intrusion policies or policy revisions are highlighted:

- Blue indicates that the highlighted setting is different in the two policies or policy revisions, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one policy or policy revision but not the other.

You can perform any of the actions described in the following table.

Table 26-4 Intrusion Policy Comparison View Actions

To...	You can...
navigate individually through changes	click **Previous** or **Next** above the title bar. The double-arrow icon ( ) centered between the left and right sides moves, and the **Difference** number adjusts to identify which difference you are viewing.
generate a new intrusion policy comparison view	click **New Comparison.** The **Select Comparison** window appears. See Using the Intrusion Policy Comparison Report for more information.
generate an intrusion policy comparison report	click **Comparison Report**. The policy comparison report creates a PDF that lists only the differences between the two policies or policy revisions.

Using the Intrusion Policy Comparison Report

License: Protection

An intrusion policy comparison report is a record of all differences between two intrusion policies or two revisions of the same intrusion policy identified by the intrusion policy comparison view, presented as a PDF. You can use this report to further examine the differences between two intrusion policy configurations and to save and disseminate your findings.

You can generate an intrusion policy comparison report from the comparison view for any intrusion policies to which you have access. Remember to commit any potential changes before you generate an intrusion policy report; only committed changes appear in the report.

The format of the intrusion policy comparison report is the same as the intrusion policy report with one exception: the intrusion policy report contains all settings in the intrusion policy, and the intrusion policy comparison report lists only those settings which differ between the policies.

Depending on your configuration, an intrusion policy comparison report can contain one or more sections as described in the Intrusion Policy Report Sections table.

Tip You can use a similar procedure to compare SSL, access control, network analysis, file, or system policies.

To compare two intrusion policies or two revisions of the same policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy**.

The Intrusion Policy page appears.

Step 2 Click **Compare Policies**.

The Select Comparison window appears.

Step 3 From the **Compare Against** drop-down list, select the type of comparison you want to make:

- To compare two different policies, select **Other Policy**.
- To compare two revisions of the same policy, select **Other Revision**.

Remember to commit any changes before you generate an intrusion policy report; only committed changes appear in the report.

Step 4 Depending on the comparison type you selected, you have the following choices:

- If you are comparing two different policies, select the policies you want to compare from the **Policy A** and **Policy B** drop-down lists.
- If you are comparing two revisions of the same policy, select the policy from the **Policy** drop-down list, then select the revisions you want to compare from the **Revision A** and **Revision B** drop-down lists.

Step 5 Click **OK** to display the intrusion policy comparison view.

The comparison view appears.

Step 6 Click **Comparison Report** to generate the intrusion policy comparison report.

Step 7 The intrusion policy report appears. You are prompted to save the report to your computer.



Tuning Intrusion Policies Using Rules

You can use the Rules page in an intrusion policy to configure rule states and other settings for shared object rules, standard text rules, and preprocessor rules.

You enable a rule by setting its rule state to Generate Events or to Drop and Generate Events. Enabling a rule causes the system to generate events on traffic matching the rule. Disabling a rule stops processing of the rule. Optionally, you can set your intrusion policy so that a rule set to Drop and Generate Events in an inline deployment generates events on, and drops, matching traffic. See Setting Drop Behavior in an Inline Deployment, page 26-5 for more information. In a passive deployment, a rule set to Drop and Generate Events just generates events on matching traffic.

You can filter rules to display a subset of rules, enabling you to select the exact set of rules where you want to change rule states or rule settings.

When an intrusion rule or rule argument requires a disabled preprocessor, the system automatically uses it with its current configuration even though it remains disabled in the network analysis policy's user interface. For more information, see Limitations of Custom Policies, page 18-11.

See the following sections for more information:

- Understanding Intrusion Prevention Rule Types, page 27-1 describes the intrusion rules and preprocessor rules you can view and configure in an intrusion policy.
- Viewing Rules in an Intrusion Policy, page 27-2 describes how you can change the order of rules on the Rules page, interpret the icons on the page, and focus in on rule details.
- Filtering Rules in an Intrusion Policy, page 27-9 describes how you can use rule filters to find the rules for which you want to apply rule settings.
- Setting Rule States, page 27-19 describes how to enable and disable rules from the Rules page.
- Filtering Intrusion Event Notification Per Policy, page 27-21 explains how to set event filtering thresholds for specific rules and set suppression on specific rules.
- Adding Dynamic Rule States, page 27-28 explains how to set rule states that trigger dynamically when rate anomalies are detected in matching traffic.
- Adding SNMP Alerts, page 27-31 describes how to associate SNMP alerts with specific rules.
- Adding Rule Comments, page 27-32 describes how to add comments to rules in an intrusion policy.

Understanding Intrusion Prevention Rule Types

License: Protection

An intrusion policy contains two types of rules: intrusion rules and preprocessor rules.

An intrusion rule is a specified set of keywords and arguments that detects attempts to exploit vulnerabilities on your network; an intrusion rule analyzes network traffic to check if it matches the criteria in the rule. The system compares packets against the conditions specified in each rule and, if the packet data matches all the conditions specified in a rule, the rule triggers. The system includes two types of intrusion rules created by the Cisco Vulnerability Research Team (VRT): shared object rules, which are compiled and cannot be modified (except for rule header information such as source and destination ports and IP addresses), and standard text rules, which can be saved and modified as new custom instances of the rule.

The system also includes preprocessor rules, which are rules associated with preprocessor and packet decoder detection options. You cannot copy or edit preprocessor rules. Most preprocessor rules are disabled by default and must be enabled (that is, set to Generate Events or to Drop and Generate Events) if you want the system to generate events for preprocessor rules and, in an inline deployment, drop offending packets.

The VRT determines the default rule states of Cisco's shared object rules, standard text rules, and preprocessor rules for each default intrusion policy included with the system.

The following table describes each type of rule included with the ASA FirePOWER module.

Table 27-1 Rule types

Type	Description
shared object rule	An intrusion rule created by the Cisco Vulnerability Research Team (VRT) that is delivered as a binary module compiled from C source code. You can use shared object rules to detect attacks in ways that standard text rules cannot. You cannot modify the rule keywords and arguments in a shared object rule; you are limited to either modifying variables used in the rule, or modifying aspects such as the source and destination ports and IP addresses and saving a new instance of the rule as a custom shared object rule. A shared object rule has a GID (generator ID) of 3. See Modifying Existing Rules, page 30-103 for more information.
standard text rule	An intrusion rule either created by the VRT, copied and saved as a new custom rule, created using the rule editor, or imported as a local rule that you create on a local machine and import. You cannot modify the rule keywords and arguments in a standard rule created by the VRT; you are limited to either modifying variables used in the rule, or modifying aspects such as the source and destination ports and IP addresses and saving a new instance of the rule as a custom standard text rule. See Modifying Existing Rules, page 30-103, Understanding and Writing Intrusion Rules, page 30-1 and Importing Local Rule Files, page 46-14 for more information. A standard text rule created by the VRT has a GID (generator ID) of 1. Custom standard text rules that you create using the rule editor or import as local rules have a SID (Signature ID) of 1000000 or greater.
preprocessor rule	A rule associated with a detection option of the packet decoder or with one of the preprocessors included with the ASA FirePOWER module. You must enable preprocessor rules if you want them to generate events. These rules have a decoder- or preprocessor-specific GID (generator ID).

Viewing Rules in an Intrusion Policy

License: Protection

You can adjust how rules are displayed in the intrusion policy, and can sort rules by several criteria. You can also display the details for a specific rule to see rule settings, rule documentation, and other rule specifics.

The Rules page has four primary areas of functionality:

- the filtering features—for more information, see Filtering Rules in an Intrusion Policy, page 27-9

- the rule attribute menus—for more information, see Setting Rule States, page 27-19, Filtering Intrusion Event Notification Per Policy, page 27-21, Adding Dynamic Rule States, page 27-28, Adding SNMP Alerts, page 27-31, and Adding Rule Comments, page 27-32
- the rules listing—for more information, see the Rules Page Columns table.
- the rule details—for more information, see Viewing Rule Details, page 27-4

You can also sort rules by different criteria; for more information, see Sorting the Rule Display, page 27-4.

Note that the icons used as column headers correspond to the menus in the menu bar, where you access those configuration items. For example, the Rule State menu is marked with the same icon ( ) as the Rule State column.

The following table describes the columns on the Rules page.

Table 27-2 Rules Page Columns

Heading	Description	For more information, see...
GID	Integer which indicates the Generator ID (GID) for the rule.	Viewing Events, page 37-1
SID	Integer which indicates the Snort ID (SID), which acts a unique identifier for the rule.	Viewing Events, page 37-1
Message	Message included in events generated by this rule, which also acts as the name of the rule.	Defining the Event Message, page 30-11
	The rule state for the rule, which may be one of three states: • drop and generate events ( ) • generate events ( ) • disable ( ) Note that you can access the Set rule state dialog box for a rule by clicking on its rule state icon.	Setting Rule States, page 27-19
	Event filter, including event thresholds and event suppression, applied to the rule.	Filtering Intrusion Event Notification Per Policy, page 27-21
	Dynamic rule state for the rule, which goes into effect if specified rate anomalies occur.	Adding Dynamic Rule States, page 27-28
	Alerts configured for the rule (currently SNMP alerts only).	Adding SNMP Alerts, page 27-31
	Comments added to the rule.	Adding Rule Comments, page 27-32

You can also use the layer drop-down list to switch to the Rules page for other layers in your policy. Note that, unless you add layers to your policy, the only editable views listed in the drop-down list are the policy Rules page and the Rules page for a policy layer that is originally named `My Changes`; note also that making changes in one of these views is the same as making the changes in the other. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information. The drop-down list also lists the Rules page for the read-only base policy. See Understanding the Base Layer, page 19-2 for information on the base policy.

To view the rules in an intrusion policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Rules** on the Policy Information page.

The Rules page appears. By default, the page lists the rules alphabetically by message.

Note that selecting **Rules** above the dividing line in the navigation panel takes you to the same rules listing. You can view and set all rule attributes in your policy in this view.

Sorting the Rule Display

License: Protection

You can sort rules by any of the columns in the Rules page by clicking on the heading title or icon.

Note that an up (▲) or down (▶) arrow on a heading or icon indicates that the sort is on that column in that direction.

To sort rules in an intrusion policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Rules**.

The Rules page appears. By default, the page lists the rules alphabetically by message.

Step 4 Click the title or icon in the top of the column you want to sort by.

The rules are sorted by the column, in the direction indicated by the arrow that appears on the column heading. To sort in the opposite direction, click the heading again. The sort order and the arrow reverse.

Viewing Rule Details

License: Protection

You can view rule documentation and rule overhead from the Rule Detail view. You can also view and add rule-specific features.

Note that local rules do not have any overhead, unless they are mapped to a vulnerability.

Table 27-3 Rule Details

Item	Description	For more information, see...
Summary	The rule summary. For rule-based events, this row appears when the rule documentation contains summary information.	Viewing Events, page 37-1
Rule State	The current rule state for the rule. Also indicates the layer where the rule state is set.	Setting Rule States, page 27-19; Using Layers in a Network Analysis or Intrusion Policy, page 19-1
Thresholds	Thresholds currently set for this rule, as well as the facility to add a threshold for the rule.	Setting a Threshold for a Rule, page 27-6
Suppressions	Suppression settings currently set for this rule, as well as the facility to add suppressions for the rule.	Setting Suppression for a Rule, page 27-6
Dynamic State	Rate-based rule states currently set for this rule, as well as the facility to add dynamic rule states for the rule.	Setting a Dynamic Rule State for a Rule, page 27-7
Alerts	Alerts currently set for this rule, as well as the facility to add an alert for the rule. Currently, only SNMP alerts are supported.	Setting an SNMP Alert for a Rule, page 27-8
Comments	Comments added to this rule, as well as the facility to add comments for the rule.	Adding a Rule Comment for a Rule, page 27-8
Documentation	The rule documentation for the current rule, supplied by the Cisco Vulnerability Research Team (VRT).	Viewing Events, page 37-1

To view rule details:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Rules**.

The Rules page appears. By default, the page lists the rules alphabetically by message.

Step 4 Highlight the rule whose rule details you want to view.

Step 5 Click **Show details**.

The Rule Detail view appears. To hide the details again, click **Hide details**.

Tip You can also open Rule Detail by double-clicking a rule in the Rules view.

Setting a Threshold for a Rule

License: Protection

You can set a single threshold for a rule from the Rule Detail page. Adding a threshold overwrites any existing threshold for the rule. For more information on thresholding, see Configuring Event Thresholding, page 27-21.

Note that a revert icon ( ) appears in a field when you type an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

To set a threshold from the rule details:

Step 1 Click **Add** next to **Thresholds**.

The Set Threshold dialog box appears.

Step 2 From the **Type** drop-down list, select the type of threshold you want to set:

- Select **Limit** to limit notification to the specified number of event instances per time period.
- Select **Threshold** to provide notification for each specified number of event instances per time period.
- Select **Both** to provide notification once per time period after a specified number of event instances.

Step 3 From the **Track By** drop-down list, select **Source** or **Destination** to indicate whether you want the event instances tracked by source or destination IP address.

Step 4 In the **Count** field, type the number of event instances you want to use as your threshold.

Step 5 In the **Seconds** field, type a number between 0 and 2147483647 that specifies the time period, in seconds, for which event instances are tracked.

Step 6 Click **OK**.

The system adds your threshold and displays an event filter icon ( ) next to the rule in the Event Filtering column. If you add multiple event filters to a rule, the system includes an indication over the icon of the number of event filters.

Setting Suppression for a Rule

License: Protection

You can set one or more suppressions for a rule from the Rule Detail page. For more information on suppression, see Configuring Suppression Per Intrusion Policy, page 27-25.

Note that a revert icon ( ) appears in a field when you type an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

To set suppression from the rule details:

Step 1 Click **Add** next to **Suppressions**.

The Add Suppression dialog box appears.

Step 2 From the **Suppression Type** drop-down list, select one of the following options:

- Select **Rule** to completely suppress events for a selected rule.
- Select **Source** to suppress events generated by packets originating from a specified source IP address.
- Select **Destination** to suppress events generated by packets going to a specified destination IP address.

Step 3 If you selected **Source** or **Destination** for the suppression type, the Network field appears. In the **Network** field, enter the IP address, an address block, or a comma-separated list comprised of any combination of these. If the intrusion policy is associated with the default action of an access control policy, you can also specify or list a network variable in the default action variable set.

For information on using IPv4 CIDR and IPv6 prefix length address blocks, see IP Address Conventions, page 1-4.

Step 4 Click **OK**.

The system adds your suppression conditions and displays an event filter icon ( ) next to the rule in the Event Filtering column next the suppressed rule. If you add multiple event filters to a rule, a number over the icon indicates the number of filters.

Setting a Dynamic Rule State for a Rule

License: Protection

You can set one or more dynamic rule states for a rule from the Rule Detail page. The first dynamic rule state listed has the highest priority. Note that when two dynamic rule states conflict, the action of the first is carried out. For more information on dynamic rule states, see Understanding Dynamic Rule States, page 27-28.

Note that a revert icon ( ) appears in a field when you type an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

To set a dynamic rule state from the rule details:

Step 1 Click **Add** next to **Dynamic State**.

The Add Rate-Based Rule State dialog box appears.

Step 2 From the **Track By** drop-down list, select an option to indicate how you want the rule matches tracked:

- Select **Source** to track the number of hits for that rule from a specific source or set of sources.
- Select **Destination** to track the number of hits for that rule to a specific destination or set of destinations.
- Select **Rule** to track all matches for that rule.

Step 3 Optionally, if you set **Track By** to **Source** or **Destination**, enter the IP address of each host you want to track in the **Network** field.

For information on using IPv4 CIDR and IPv6 prefix length notation, see IP Address Conventions, page 1-4.

Step 4 Next to **Rate**, indicate the number of rule matches per time period to set the attack rate:

- In the **Count** field, using an integer between 0 and 2147483647, specify the number of rule matches you want to use as your threshold.
- In the **Seconds** field, using an integer between 0 and 2147483647, specify the number of seconds that make up the time period for which attacks are tracked.

Step 5 From the **New State** drop-down list, select the new action to be taken when the conditions are met:

- Select **Generate Events** to generate an event.
- Select **Drop and Generate Events** to generate an event and drop the packet that triggered the event in inline deployments or to generate an event in passive deployments.
- Select **Disabled** to take no action.

Step 6 In the **Timeout** field, using an integer between 1 and 2147483647 (approximately 68 years), type the number of seconds you want the new action to remain in effect. After the timeout occurs, the rule reverts to its original state. Specify `0` to prevent the new action from timing out.

Step 7 Click **OK**.

The system adds the dynamic rule state and displays a dynamic state icon ( ) next to the rule in the Dynamic State column. If you add multiple dynamic rule state filters to a rule, a number over the icon indicates the number of filters.

If any required fields are blank, you receive an error message indicating which fields you must fill.

Setting an SNMP Alert for a Rule

License: Protection

You can set an SNMP alert for a rule from the Rule Detail page. For more information on SNMP alerts, see Adding SNMP Alerts, page 27-31.

To add an SNMP alert from the rule details:

Step 1 Click **Add SNMP Alert** next to **Alerts**.

The system adds the alert and displays an alert icon ( ) next to the rule in the Alerting column. If you add multiple alerts to a rule, the system includes an indication over the icon of the number of alerts.

Adding a Rule Comment for a Rule

License: Protection

You can add a rule comment for a rule from the Rule Detail page. For more information on rule comments, see Adding Rule Comments, page 27-32.

To add a comment from the rule details:

Step 1 Click **Add** next to **Comments**.

The Add Comment dialog box appears.

Step 2 In the **Comment** field, type the rule comment.

Step 3 Click **OK**.

The system adds the comment and displays a comment icon ( ) next to the rule in the Comments column. If you add multiple comments to a rule, a number over the icon indicates the number of comments.

Tip To delete a rule comment, click **Delete** in the rule comments section. Note that you can only delete a comment if the comment is cached with uncommitted intrusion policy changes. After intrusion policy changes are committed, the rule comment is permanent.

Filtering Rules in an Intrusion Policy

License: Protection

You can filter the rules you display on the Rules page by a single criteria, or a combination of one or more criteria.

The filter you construct is shown in the Filter text box. You can click keywords and keyword arguments in the filter panel to construct a filter. When you select multiple keywords, the system combines them using AND logic to create a compound search filter. For example, if you select **preprocessor** under **Category** and then select **Rule Content > GID** and enter `116`, you get a filter of `Category: "preprocessor" GID:"116"` which retrieves all rules that are preprocessor rules **and** have a GID of 116.

The Category, Microsoft Vulnerabilities, Microsoft Worms, Platform Specific, Preprocessor, and Priority filter groups allow you to submit more than one argument for a keyword, separated by commas. For example, you can press Shift and then select **os-linux** and **os-windows** from **Category** to produce the filter `Category:"os-windows,os-linux"`, which retrieves any rules in the `os-linux` category or in the `os-windows` category.

To show the filter panel, click the show icon ( ).

To hide the filter panel, click the hide icon ( ).

For more information, see the following topics:

- Understanding Rule Filtering in an Intrusion Policy, page 27-9
- Setting a Rule Filter in an Intrusion Policy, page 27-17

Understanding Rule Filtering in an Intrusion Policy

License: Protection

Rule filter keywords help you find the rules for which you want to apply rule settings, such as rule states or event filters. You can filter by a keyword and simultaneously select the argument for the keyword by selecting the argument you want from the Rules page filter panel.

For more information, see the following sections:

- Guidelines for Constructing Intrusion Policy Rule Filters, page 27-10

- Understanding Rule Configuration Filters, page 27-12
- Understanding Rule Content Filters, page 27-14
- Understanding Rule Categories, page 27-16
- Editing a Rule Filter Directly, page 27-16

Guidelines for Constructing Intrusion Policy Rule Filters

License: Protection

In most cases, when you are building a filter, you can use the filter panel to the left of the Rules page in the intrusion policy to select the keywords/arguments you want to use.

Rule filters are grouped into rule filter groups in the filter panel. Many rule filter groups contain sub-criteria so that you can more easily find the specific rules you are looking for. Some rule filters have multiple levels that you can expand to drill down to individual rules.

Items in the filter panel sometimes represent filter type groups, sometimes represent keywords, and sometimes represent the argument to a keyword. Use the following rules of thumb to help you build your filters:

- When you select a filter type group heading that is not a keyword (Rule Configuration, Rule Content, Platform Specific, and Priority), it expands to list the available keywords.

  When you select a keyword by clicking on a node in the criteria list, a pop-up window appears, where you supply the argument you want to filter by.

  If that keyword is already used in the filter, the argument you supply replaces the existing argument for that keyword.

- When you select a filter type group heading that is a keyword (Category, Classifications, Microsoft Vulnerabilities, Microsoft Worms, Priority, and Rule Update), it lists the available arguments.

  When you select an item from this type of group, the argument and the keyword it applies to are immediately added to the filter. If the keyword is already in the filter, it replaces the existing argument for the keyword that corresponds to that group.

  For example, if you click **os-linux** under **Category** in the filter panel, `Category:"os-linux"` is added to the filter text box. If you then click **os-windows** under **Category**, the filter changes to `Category:"os-windows"`.

- Reference under Rule Content is a keyword, and so are the specific reference ID types listed below it. When you select any of the reference keywords, a pop-up window appears, where you supply an argument and the keyword is added to the existing filter. If the keyword is already in use in the filter, the new argument you supply replaces the existing argument.

  For example, if you click **Rule Content > Reference > CVE ID** in the filter panel, a pop-up window prompts you to supply the CVE ID. If you enter `2007`, then `CVE:"2007"` is added to the filter text box. In another example, if you click **Rule Content > Reference** in the filter panel, a pop-up window prompts you to supply the reference. If you enter `2007`, then `Reference:"2007"` is added to the filter text box.

- When you select rule filter keywords from different groups, each filter keyword is added to the filter and any existing keywords are maintained (unless overridden by a new value for the same keyword).

  For example, if you click **os-linux** under **Category** in the filter panel, `Category:"os-linux"` is added to the filter text box. If you then click **MS00-006** under **Microsoft Vulnerabilities**, the filter changes to `Category:"os-linux" MicrosoftVulnerabilities:"MS00-006"`.

- When you select multiple keywords, the system combines them using AND logic to create a compound search filter. For example, if you select **preprocessor** under **Category** and then select **Rule Content > GID** and enter `116`, you get a filter of `Category: "preprocessor" GID:"116"`, which retrieves all rules that are preprocessor rules **and** have a GID of 116.
- The Category, Microsoft Vulnerabilities, Microsoft Worms, Platform Specific, and Priority filter groups allow you to submit more than one argument for a keyword, separated by commas. For example, you can press Shift and then select **os-linux** and **os-windows** from **Category** to produce the filter `Category:"os-windows,app-detect"`, which retrieves any rules in the `os-linux` category or in the `os-windows` category.

The same rule may be retrieved by more than one filter keyword/argument pair. For example, the DOS Cisco attempt rule (SID 1545) appears if rules are filtered by the **dos** category, and also if you filter by the **High** priority.

Note The Cisco VRT may use the rule update mechanism to add and remove rule filters.

Note that the rules on the Rules page may be either shared object rules (generator ID 3) or standard text rules (generator ID 1). The following table describes the different rule filters.

Table 27-4 Rule Filter Groups

Filter Group	Description	Multiple Argument Support?	Heading is...	Items in List are...
Rule Configuration	Finds rules according to the configuration of the rule. See Understanding Rule Configuration Filters, page 27-12.	No	A grouping	keywords
Rule Content	Finds rules according to the content of the rule. See Understanding Rule Content Filters, page 27-14.	No	A grouping	keywords
Category	Finds rules according to the rule categories used by the rule editor. Note that local rules appear in the local sub-group. See Understanding Rule Categories, page 27-16.	Yes	A keyword	arguments
Classifications	Finds rules according to the attack classification that appears in the packet display of an event generated by the rule. See Defining the Intrusion Event Classification, page 30-12.	No	A keyword	arguments
Microsoft Vulnerabilities	Finds rules according to Microsoft bulletin number.	Yes	A keyword	arguments
Microsoft Worms	Finds rules based on specific worms that affect Microsoft Windows hosts.	Yes	A keyword	arguments
Platform Specific	Finds rules according to their relevance to specific versions of operating systems. Note that a rule may affect more than one operating system or more than one version of an operating system. For example, enabling SID 2260 affects multiple versions of Mac OS X, IBM AIX, and other operating systems.	Yes	A keyword	arguments Note that if you pick one of the items from the sub-list, it adds a modifier to the argument.

Table 27-4 Rule Filter Groups (continued)

Filter Group	Description	Multiple Argument Support?	Heading is...	Items in List are...
Preprocessors	Finds rules for individual preprocessors. Note that you must enable preprocessor rules associated with a preprocessor option to generate events for the option when the preprocessor is enabled; see Setting Rule States, page 27-19.	Yes	A grouping	sub-groupings
Priority	Finds rules according to high, medium, and low priorities. The classification assigned to a rule determines its priority. These groups are further grouped into rule categories. Note that local rules (that is, rules that you create) do not appear in the priority groups.	Yes	A keyword	arguments Note that if you pick one of the items from the sub-list, it adds a modifier to the argument.
Rule Update	Finds rules added or modified through a specific rule update. For each rule update, view all rules in the update, only new rules imported in the update, or only existing rules changed by the update.	No	A keyword	arguments

Understanding Rule Configuration Filters

License: Protection

You can filter the rules listed in the Rules page by several rule configuration settings.

When you select a keyword by clicking on a node in the criteria list, a pop-up window appears, where you supply the argument you want to filter by.

If that keyword is already used in the filter, the argument you supply replaces the existing argument for that keyword.

See the following procedures for more information on the rule configuration settings you can use to filter.

To use the Rule State filter:

Step 1 Under **Rule Configuration**, click **Rule State**.

Step 2 From the **Rule State** drop-down list, select the rule state to filter by:

- To find rules that only generate events, select **Generate Events**, then click **OK**.
- To find rules that are set to generate events and drop the matching packet, select **Drop and Generate Events**, then click **OK**.
- To find disabled rules, select **Disabled**, then click **OK**.

The Rules page updates to display rules according to current rule state.

To use the Threshold filter:

Step 1 Under **Rule Configuration**, click **Threshold**.

Step 2 From the **Threshold** drop-down list, select the threshold setting to filter by:

- To find rules with a threshold type of `limit`, select **Limit**, then click **OK**.
- To find rules with a threshold type of `threshold`, select **Threshold**, then click **OK**.
- To find rules with a threshold type of `both`, select **Both**, then click **OK**.
- To find rules with thresholds tracked by `source`, select **Source**, then click **OK**.
- To find rules with thresholds tracked by `destination`, select **Destination**, then click **OK**.
- To find any rule with a threshold set, select **All**, then click **OK**.

The Rules page updates to display rules where the type of threshold indicated in the filter has been applied to the rule.

To use the Suppression filter:

Step 1 Under **Rule Configuration**, click **Suppression**.

Step 2 From the **Suppression** drop-down list, select the suppression setting to filter by:

- To find rules where events are suppressed for packets inspected by that rule, select **By Rule**, then click **OK**.
- To find rules where events are suppressed based on the source of the traffic, select **By Source**, then click **OK**.
- To find rules where events are suppressed based on the destination of the traffic, select **By Destination**, then click **OK**.
- To find any rule with suppression set, select **All**, then click **OK**.

The Rules page updates to display rules where the type of suppression indicated in the filter has been applied to the rule.

To use the Dynamic State filter:

Step 1 Under **Rule Configuration**, click **Dynamic State**.

Step 2 From the **Dynamic State** drop-down list, select the suppression setting to filter by:

- To find rules where a dynamic state is configured for packets inspected by that rule, select **By Rule**, then click **OK**.
- To find rules where a dynamic state is configured for packets based on the source of the traffic, select **By Source**, then click **OK**.
- To find rules where a dynamic state is configured based on the destination of the traffic, select **By Destination**, then click **OK**.
- To find rules where a dynamic state of `Generate Events` is configured, select **Generate Events**, then click **OK**.

- To find rules where a dynamic state of `Drop and Generate Events` is configured, select **Drop and Generate Events**, then click **OK**.
- To find where a dynamic state of `Disabled` is configured, select **Disabled**, then click **OK**.
- To find any rule with suppression set, select **All**, then click **OK**.

The Rules page updates to display rules where the dynamic rule state indicated in the filter has been applied to the rule.

To use the Alert filter:

Step 1 Under **Rule Configuration**, click **Alert**.

Step 2 From the **Alert** drop-down list, select the alert setting to filter by: **SNMP**.

Step 3 Click **OK**.

The Rules page updates to display rules where you have applied an alert filter.

To use the Comment filter:

Step 1 Under **Rule Configuration**, click **Comment**.

Step 2 In the **Comment** field, type the string of comment text to filter by, then click **OK**.

The Rules page updates to display rules where comments applied to the rule contain the string indicated in the filter.

Understanding Rule Content Filters

License: Protection

You can filter the rules listed in the Rules page by several rule content items. For example, you can quickly retrieve a rule by searching for the rule's SID. You can also find all rules that inspect traffic going to a specific destination port.

When you select a keyword by clicking on a node in the criteria list, a pop-up window appears, where you supply the argument you want to filter by.

If that keyword is already used in the filter, the argument you supply replaces the existing argument for that keyword.

For example, if you click **SID** under **Rule Content** in the filter panel, a pop-up window appears, prompting you to supply a SID. If you type `1045`, then `SID:"1045"`is added to the filter text box. If you then click **SID** again and change the SID filter to `1044`, the filter changes to `SID:"1044"`.

For more information on the rule content you can use to filter, see the following table.

Table 27-5 Rule Content Filters

To use this filter, click...	Then...	Result
Message	Type the message string to filter by, then click **OK**.	Finds rules that contain the supplied string in the message field.
SID	Type the SID number to filter by, then click **OK**.	Finds rules that have the specified SID.
GID	Type the GID number to filter by, then click **OK**.	Finds rules that have the specified GID.
Reference	Type the reference string to filter by, then click **OK**. To enter a string for a specific type of reference that you want to filter by, select **CVE ID**, **URL**, **Bugtraq ID**, **Nessus ID**, **Arachnids ID**, or **Mcafee ID**, then type a string and click **OK**.	Finds rules that contain the supplied string in the reference field.
Action	Select the action to filter by: • To find alert rules, select **Alert**, then click **OK**. • To find pass rules, select **Pass**, then click **OK**.	Finds rules that start with `alert` or `pass`.
Protocol	Select the protocol to filter by: **ICMP**, **IP**, **TCP**, or **UDP**; then click **OK**.	Finds rules that include the selected protocol.
Direction	Select a directional setting to filter by: • To find rules that inspect traffic moving in a specific direction, select **Directional**, then click **OK**. • To find rules that inspect traffic moving in either direction between a source and destination, select **Bidirectional**, then click **OK**.	Finds rules based on whether the rule includes the indicated directional setting.
Source IP	Type the source IP address to filter by, then click **OK**. Note that you can filter by a valid IP address, a CIDR block/prefix length, or using variables such as `$HOME_NET` or `$EXTERNAL_NET`.	Finds rules that use the specified addresses or variables for the source IP address designation in the rule.
Destination IP	Type the destination IP address to filter by, then click **OK**. Note that you can filter by a valid IP address, a CIDR block/prefix length, or using variables such as `$HOME_NET` or `$EXTERNAL_NET`.	Finds rules that use the specified addresses or variables for the source IP address designation in the rule.

Table 27-5 Rule Content Filters (continued)

To use this filter, click...	Then...	Result
Source port	Type the source port to filter by, then click **OK**. The port value must be an integer between 1 and 65535 or a port variable.	Finds rules that include the specified source port.
Destination port	Type the destination port to filter by, then click **OK**. The port value must be an integer between 1 and 65535 or a port variable.	Finds rules that include the specified destination port.
Rule Overhead	Select the amount of rule overhead to filter by: **Low**, **Medium**, **High**, or **Very High**; then click **OK**.	Finds rules with the selected rule overhead.
Metadata	Type the metadata key-value pair to filter by, separated by a space; then click **OK**. For example, type `metadata:"service http"` to locate rules with metadata relating to the HTTP application protocol.	Find rules with metadata containing the matching key-value pair.

Understanding Rule Categories

License: Protection

The ASA FirePOWER module places rules in categories based on the type of traffic the rule detects. On the Rules page, you can filter by rule category, so you can set a rule attribute for all rules in a category. For example, if you do not have Linux hosts on your network, you could filter by the **os-linux** category, then disable all the rules showing to disable the entire **os-linux** category.

Note The Cisco VRT may use the rule update mechanism to add and remove rule categories.

Editing a Rule Filter Directly

License: Protection

You can edit your filter to modify the special keywords and their arguments that are supplied when you click on a filter in the filter panel. Custom filters on the Rules page function like those used in the rule editor, but you can also use any of the keywords supplied in the Rules page filter, using the syntax displayed when you select the filter through the filter panel. To determine a keyword for future use, click on the appropriate argument in the filter panel on the right. The filter keyword and argument syntax appear in the filter text box.

To see lists of arguments for keywords which only support specific values, see Understanding Rule Configuration Filters, page 27-12, Understanding Rule Content Filters, page 27-14, and Understanding Rule Categories, page 27-16. Remember that comma-separated multiple arguments for a keyword are only supported for the Category and Priority filter types.

You can use keywords and arguments, character strings, and literal character strings in quotes, with spaces separating multiple filter conditions. A filter cannot include regular expressions, wild card characters, or any special operator such as a negation character (!), a greater than symbol (>), less than symbol (<), and so on. When you type in search terms without a keyword, without initial capitalization of the keyword, or without quotes around the argument, the search is treated as a string search and the category, message, and SID fields are searched for the specified terms.

All keywords, keyword arguments, and character strings are case-insensitive. Except for the `gid` and `sid` keywords, all arguments and strings are treated as partial strings. Arguments for `gid` and `sid` return only exact matches.

Each rule filter can include one or more keywords in the format:

`Keyword:"argument"`

where *`keyword`* is one of the keywords in the filter groups described in the Rule types table and *`argument`* is enclosed in double quotes and is a single, case-insensitive, alphanumeric string to search for in the specific field or fields relevant to the keyword. Note that keywords should be typed with initial capitalization.

Arguments for all keywords except `gid` and `sid` are treated as partial strings. For example, the argument `123` returns `"12345"`, `"41235"`, `"45123"`, and so on. The arguments for `gid` and `sid` return only exact matches; for example, `sid:3080` returns only `SID 3080`.

Each rule filter can also include one or more alphanumeric character strings. Character strings search the rule Message field, Signature ID, and Generator ID. For example, the string `123` returns the strings `"Lotus123"`, `"123mania"`, and so on in the rule message, and also returns `SID 6123`, `SID 12375`, and so on. For information on the rule Message field, see Defining the Event Message, page 30-11. You can search for a partial SID by filtering with one or more character strings.

All character strings are case-insensitive and are treated as partial strings. For example, any of the strings `ADMIN`, `admin`, or `Admin` return `"admin"`, `"CFADMIN"`, `"Administrator"` and so on.

You can enclose character strings in quotes to return exact matches. For example, the literal string `"overflow attempt"` in quotes returns only that exact string, whereas a filter comprised of the two strings `overflow` and `attempt` without quotes returns `"overflow attempt"`, `"overflow multipacket attempt"`, `"overflow with evasion attempt"`, and so on.

You can narrow filter results by entering any combination of keywords, character strings, or both, separated by spaces. The result includes any rule that matches all the filter conditions.

You can enter multiple filter conditions in any order. For example, each of the following filters returns the same rules:

- `url:at login attempt cve:200`
- `login attempt cve:200 url:at`
- `login cve:200 attempt url:at`

Setting a Rule Filter in an Intrusion Policy

License: Protection

You can filter the rules on the Rules page to display a subset of rules. You can then use any of the page features. This can be useful, for example, when you want to set a threshold for all the rules in a specific category. You can use the same features with rules in a filtered or unfiltered list. For example, you can apply new rule states to rules in a filtered or unfiltered list.

You can select predefined filter keywords from the filter panel on the left side of the Rules page in the intrusion policy. When you select a filter, the page displays all matching rules, or indicates when no rules match.

For more information on all the keywords and arguments you can use and how you can construct filters from the filter panel, see Understanding Rule Filtering in an Intrusion Policy, page 27-9.

You can add keywords to a filter to further constrain it. Any filter you enter searches the entire rules database and returns all matching rules. When you enter a filter while the page still displays the result of a previous filter, the page clears and returns the result of the new filter instead.

You can also type a filter using the same keyword and argument syntax supplied when you select a filter, or modify argument values in a filter after you select it. When you type in search terms without a keyword, without initial capitalization of the keyword, or without quotes around the argument, the search is treated as a string search and the category, message, and SID fields are searched for the specified terms.

To filter for specific rules in an intrusion policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Rules**.

The Rules page appears. By default, the page lists the rules alphabetically by message.

Step 4 Construct a filter by clicking on keywords or arguments in the filter panel on the left. Note that if you click an argument for a keyword already in the filter, it replaces the existing argument. See the following for more information:

- Guidelines for Constructing Intrusion Policy Rule Filters, page 27-10
- Understanding Rule Configuration Filters, page 27-12
- Understanding Rule Content Filters, page 27-14
- Understanding Rule Categories, page 27-16
- Editing a Rule Filter Directly, page 27-16

The page refreshes to display all matching rules, and the number of rules matching the filter is displayed above the filter text box.

Step 5 Select the rule or rules where you want to apply a new setting. You have the following options:

- To select a specific rule, select the check box next to the rule.
- To select all the rules in the current list, select the check box at the top of the column.

Step 6 Optionally, make any changes to the rule that you would normally make on the page. See the following sections for more information:

- See Setting Rule States, page 27-19 for information on enabling and disabling rules on the Rules page.
- See Filtering Intrusion Event Notification Per Policy, page 27-21 for information on adding thresholding and suppression to rules.
- See Adding Dynamic Rule States, page 27-28 for information on setting dynamic rule states that trigger when rate anomalies occur in matching traffic.

- See Adding SNMP Alerts, page 27-31 for information on adding SNMP alerts to specific rules.
- See Adding Rule Comments, page 27-32 for information on adding rule comments to rules.

Step 7 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache.

See Managing Intrusion Policies, page 26-3 and Editing Intrusion Policies, page 26-4 for more information.

Setting Rule States

License: Protection

The Cisco Vulnerability Research Team (VRT) sets the default state of each intrusion and preprocessor rule in each default policy. For example, a rule may be enabled in the Security over Connectivity default policy and disabled in the Connectivity over Security default policy. Intrusion policy rules you create inherit the default states of the rules in the default policy you use to create your policy.

You can set a rule to Generate Events, to Drop and Generate Events, or to Disable individually, or you can filter the rules by a variety of factors to select the rules for which you want to modify the state. In an inline deployment, you can use the Drop and Generate Events rule state in inline intrusion deployments to drop malicious packets. Note that rules with the Drop and Generate Events rule state generate events but do not drop packets in a passive deployment. Setting a rule to Generate Events or to Drop and Generate Events enables the rule; setting the rule to Disable disables it.

Consider two scenarios. In the first scenario, the rule state for a specific rule is set to Generate Events. When a malicious packet crosses your network and triggers the rule, the packet is sent to its destination and the system generates an intrusion event. In the second scenario, assume that the rule state for the same rule is set to Drop and Generate Events in an inline deployment. In this case, when the malicious packet crosses the network, the system drops the malicious packet and generates an intrusion event. The packet never reaches its target.

In an intrusion policy, you can set a rule's state to one of the following:

- Set the rule state to **Generate Events** if you want the system to detect a specific intrusion attempt and generate an intrusion event when it finds matching traffic.
- Set the rule state to **Drop and Generate Events** if you want the system to detect a specific intrusion attempt, then drop the packet containing the attack and generate an intrusion event when it finds matching traffic in an inline deployment, or to generate an intrusion event when it finds matching traffic in a passive deployment.

  Note that for the system to drop packets, your intrusion policy must be set to drop rules in an inline deployment; see Setting Drop Behavior in an Inline Deployment, page 26-5 for more information.
- Set the rule state to **Disable** if you do not want the system to evaluate matching traffic.

To use drop rules, you must:

- Enable the **Drop when Inline** option in your intrusion policy.
- Set the rule state to **Drop and Generate Events** for any rules that should drop all packets that match the rule.
- Apply an access control policy that includes an access control rule that is associated with your intrusion policy in an inline deployment.

Filtering rules on the Rules page can help you find the rules you want to set as drop rules. For more information, see Filtering Rules in an Intrusion Policy, page 27-9.

See Understanding and Writing Intrusion Rules, page 30-1 for information about rule anatomy, rule keywords and their options, and rule writing syntax.

The VRT sometimes uses a rule update to change the default state of one or more rules in a default policy. If you allow rule updates to update your base policy, you also allow the rule update to change the default state of a rule in your policy when the default state changes in the default policy you used to create your policy (or in the default policy it is based on). Note, however, that if you have changed the rule state, the rule update does not override your change.

To change the rule state for one or more rules:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Note that this page indicates the total number of enabled rules, the total number of enabled rules set to Generate Events, and the total number set to Drop and Generate Events. Note also that in a passive deployment, rules set to Drop and Generate Events only generate events.

Step 3 Click **Rules**.

The Rules page appears. By default, the page lists the rules alphabetically by message.

Step 4 Locate the rule or rules where you want to set the rule state. You have the following options:

- To sort the current display, click on a column heading or icon. To reverse the sort, click again.
- Construct a filter by clicking on keywords or arguments in the filter panel on the left. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy, page 27-9 and Setting a Rule Filter in an Intrusion Policy, page 27-17.

  The page refreshes to display all matching rules.

Step 5 Select the rule or rules where you want to set the rule state. You have the following options:

- To select a specific rule, select the check box next to the rule.
- To select all the rules in the current list, select the check box at the top of the column.

Step 6 You have the following options:

- To generate events when traffic matches the selected rules, select **Rule State > Generate Events**.
- To generate events and drop the traffic in inline deployments when traffic matches the selected rules, select **Rule State > Drop and Generate Events**.
- To not inspect traffic matching the selected rules, select **Rule State > Disable**.

Note Cisco **strongly** recommends that you **do not** enable all the intrusion rules in an intrusion policy. The performance of your device is likely to degrade if all rules are enabled. Instead, tune your rule set to match your network environment as closely as possible.

Step 7 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See Managing Intrusion Policies, page 26-3 and Editing Intrusion Policies, page 26-4 for more information.

Filtering Intrusion Event Notification Per Policy

License: Protection

The importance of an intrusion event can be based on frequency of occurrence, or on source or destination IP address. In some cases you may not care about an event until it has occurred a certain number of times. For example, you may not be concerned if someone attempts to log into a server until they fail a certain number of times. In other cases, you may only need to see a few occurrences to know there is a widespread problem. For example, if a DoS attack is launched against your web server, you may only need to see a few occurrences of an intrusion event to know that you need to address the situation. Seeing hundreds of the same event only overwhelms your system.

See the following sections for more information:

- Configuring Event Thresholding, page 27-21 explains how to set thresholds that dictate how often an event is displayed, based on the number of occurrences. You can configure thresholding per event and per policy.
- Configuring Suppression Per Intrusion Policy, page 27-25 explains how to suppress notification of specified events per source or destination IP address, per policy.

Configuring Event Thresholding

License: Protection

You can set thresholds for individual rules, per intrusion policy, to limit the number of times the system logs and displays an intrusion event based on how many times the event is generated within a specified time period. This can prevent you from being overwhelmed with a large number of identical events. You can set thresholds per shared object rule, standard text rule, or preprocessor rule.

For more information, see the following sections:

- Understanding Event Thresholding, page 27-21
- Adding and Modifying Intrusion Event Thresholds, page 27-23
- Viewing and Deleting Intrusion Event Thresholds, page 27-24
- Setting a Threshold for a Rule, page 27-6

Understanding Event Thresholding

License: Protection

First, you must specify the thresholding type. You can select from the options discussed in the following table.

Table 27-6 Thresholding Options

Option	Description
Limit	Logs and displays events for the specified number of packets (specified by the Count argument) that trigger the rule during the specified time period. For example, if you set the type to **Limit**, the **Count** to `10`, and the **Seconds** to `60`, and 14 packets trigger the rule, the system stops logging events for the rule after displaying the first 10 that occur within the same minute.
Threshold	Logs and displays a single event when the specified number of packets (specified by the Count argument) trigger the rule during the specified time period. Note that the counter for the time restarts after you hit the threshold count of events and the system logs that event. For example, you set the type to **Threshold**, **Count** to `10`, and **Seconds** to `60`, and the rule triggers 10 times by second 33. The system generates one event, then resets the Seconds and Count counters to 0. The rule then triggers another 10 times in the next 25 seconds. Because the counters reset to 0 at second 33, the system logs another event.
Both	Logs and displays an event once per specified time period, after the specified number (count) of packets trigger the rule. For example, if you set the type to **Both**, **Count** to two, and **Seconds** to `10`, the following event counts result: • If the rule is triggered once in 10 seconds, the system does not generate any events (the threshold is not met) • If the rule is triggered twice in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time) • If the rule is triggered four times in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time, and following events are ignored)

Next, you must specify tracking, which determines whether the event threshold is calculated per source or destination IP address. Select one of the options from the following table to specify how the system tracks event instances.

Table 27-7 Thresholding IP Options

Option	Description
Source	Calculates event instance count per source IP address.
Destination	Calculates event instance count per destination IP address.

Finally, you must specify the number of instances and time period that define the threshold.

Table 27-8 Thresholding Instance/Time Options

Option	Description
Count	The number of event instances per specified time period per tracking IP address required to meet the threshold.
Seconds	The number of seconds that elapse before the count resets. If you set the threshold type to **limit**, the tracking to **Source IP**, the **count** to `10`, and the **seconds** to `10`, the system logs and displays the first 10 events that occur in 10 seconds from a given source port. If only 7 events occur in the first 10 seconds, the system logs and displays those; if 40 events occur in the first 10 seconds, the system logs and displays 10, then begins counting again when the 10-second time period elapses.

Note that you can use intrusion event thresholding alone or in any combination with rate-based attack prevention, the `detection_filter` keyword, and intrusion event suppression. See Adding Dynamic Rule States, page 27-28, Filtering Events, page 30-87, and Configuring Suppression Per Intrusion Policy, page 27-25 for more information.

See the following sections for more information:

- Adding and Modifying Intrusion Event Thresholds, page 27-23
- Setting a Threshold for a Rule, page 27-6
- Viewing and Deleting Intrusion Event Thresholds, page 27-24

Adding and Modifying Intrusion Event Thresholds

License: Protection

You can set a threshold for one or more specific rules. You can also separately or simultaneously modify existing threshold settings. You can set a a single threshold for each. Adding a threshold overwrites any existing threshold for the rule.

For more information on viewing and deleting threshold configurations, see Viewing and Deleting Intrusion Event Thresholds, page 27-24.

You can also modify the global threshold that applies by default to all rules and preprocessor-generated events. For more information, see Globally Limiting Intrusion Event Logging, page 29-1.

Note that a revert icon ( ) appears in a field when you type an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

To add or modify event thresholds:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Rules**.

The Rules page appears. By default, the page lists the rules alphabetically by message.

Step 4 Locate the rule or rules where you want to set a threshold. You have the following options:

- To sort the current display, click on a column heading or icon. To reverse the sort, click again.
- Construct a filter by clicking on keywords or arguments in the filter panel on the left. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy, page 27-9 and Setting a Rule Filter in an Intrusion Policy, page 27-17.

The page refreshes to display all matching rules.

Step 5 Select the rule or rules where you want to set a threshold. You have the following options:

- To select a specific rule, select the check box next to the rule.
- To select all the rules in the current list, select the check box at the top of the column.

Step 6 Select **Event Filtering > Threshold.**

The thresholding pop-up window appears.

Step 7 From the **Type** drop-down list, select the type of threshold you want to set:

- Select **Limit** to limit notification to the specified number of event instances per time period.
- Select **Threshold** to provide notification for each specified number of event instances per time period.
- Select **Both** to provide notification once per time period after a specified number of event instances.

Step 8 From the **Track By** drop-down list, select whether you want the event instances tracked by **Source** or **Destination** IP address.

Step 9 In the **Count** field, specify the number of event instances you want to use as your threshold.

Step 10 In the **Seconds** field, specify the number of seconds that make up the time period for which event instances are tracked.

Step 11 Click **OK**.

The system adds your threshold and displays an event filter icon ( ) next to the rule in the Event Filtering column. If you add multiple event filters to a rule, a number over the icon indicates the number of event filters.

Step 12 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache.

See Managing Intrusion Policies, page 26-3 and Editing Intrusion Policies, page 26-4 for more information.

Viewing and Deleting Intrusion Event Thresholds

License: Protection

You may want to view or delete an existing threshold setting. You can use the Rules Details view to display the configured settings for a threshold to see if they are appropriate for your system. If they are not, you can add a new threshold to overwrite the existing values.

Note that you can also modify the global threshold that applies by default to all rules and preprocessor-generated events. See Globally Limiting Intrusion Event Logging, page 29-1 for more information.

To view or delete a threshold:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Rules**.

The Rules page appears. By default, the page lists the rules alphabetically by message.

Step 4 Locate the rule or rules that have a configured threshold you want to view or delete. You have the following options:

- To sort the current display, click on a column heading or icon. To reverse the sort, click again.
- Construct a filter by clicking on keywords or arguments in the filter panel on the left. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy, page 27-9 and Setting a Rule Filter in an Intrusion Policy, page 27-17.

  The page refreshes to display all matching rules.

Step 5 Select the rule or rules with a configured threshold you want to view or delete. You have the following options:

- To select a specific rule, select the check box next to the rule.
- To select all the rules in the current list, select the check box at the top of the column.

Step 6 To remove the threshold for each selected rule, select **Event Filtering > Remove Thresholds.** Click **OK** in the confirmation pop-up window that appears.

Tip To remove a specific threshold, you can also highlight the rule and click **Show details**. Expand the threshold settings, then click **Delete** next to the threshold settings you want to remove. Click **OK** to confirm that you want to delete the configuration.

The page refreshes and the threshold is deleted.

Step 7 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See Managing Intrusion Policies, page 26-3 and Editing Intrusion Policies, page 26-4 for more information.

Configuring Suppression Per Intrusion Policy

License: Protection

You can suppress intrusion event notification when a specific IP address or range of IP addresses triggers a specific rule or preprocessor. This is useful for eliminating false positives. For example, if you have a mail server that transmits packets that look like a specific exploit, you might suppress event notification for that event when it is triggered by your mail server. The rule triggers for all packets, but you only see events for legitimate attacks.

Note that you can use intrusion event suppression alone or in any combination with rate-based attack prevention, the `detection_filter` keyword, and intrusion event thresholding. See Adding Dynamic Rule States, page 27-28, Filtering Events, page 30-87, and Configuring Event Thresholding, page 27-21 for more information.

See the following sections for more information:

- Suppressing Intrusion Events, page 27-25
- Viewing and Deleting Suppression Conditions, page 27-27

Suppressing Intrusion Events

License: Protection

You can suppress intrusion event notification for a rule or rules. When notification is suppressed for a rule, the rule triggers but events are not generated. You can set one or more suppressions for a rule. The first suppression listed has the highest priority. Note that when two suppressions conflict, the action of the first is carried out.

Note that a revert icon ( ) appears in a field when you type an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

To suppress event display:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Rules**.

The Rules page appears. By default, the page lists the rules alphabetically by message.

Step 4 Locate the rule or rules where you want to set suppression. You have the following options:

- To sort the current display, click on a column heading or icon. To reverse the sort, click again.
- Construct a filter by clicking on keywords or arguments in the filter panel on the left. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy, page 27-9 and Setting a Rule Filter in an Intrusion Policy, page 27-17.

  The page refreshes to display all matching rules.

Step 5 Select the rule or rules for which you want to configure suppression conditions. You have the following options:

- To select a specific rule, select the check box next to the rule.
- To select all rules in the current list, select the check box at the top of the column.

Step 6 Select **Event Filtering > Suppression**.

The suppression pop-up window appears.

Step 7 Select one of the following **Suppression Type** options:

- Select **Rule** to completely suppress events for a selected rule.
- Select **Source** to suppress events generated by packets originating from a specified source IP address.
- Select **Destination** to suppress events generated by packets going to a specified destination IP address.

Step 8 If you selected **Source** or **Destination** for the suppression type, in the **Network** field enter the IP address, address block, or variable you want to specify as the source or destination IP address, or a comma-separated list comprised of any combination of these.

For information on using IPv4 CIDR and IPv6 prefix length address blocks, see IP Address Conventions, page 1-4.

Step 9 Click **OK**.

The system adds your suppression conditions and displays an event filter icon ( ) next to the rule in the Event Filtering column next the suppressed rule. If you add multiple event filters to a rule, a number over the icon indicates the number of event filters.

Step 10 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache.

See Managing Intrusion Policies, page 26-3 and Editing Intrusion Policies, page 26-4 for more information.

Viewing and Deleting Suppression Conditions

License: Protection

You may want to view or delete an existing suppression condition. For example, you can suppress event notification for packets originating from a mail server IP address because the mail server normally transmits packets that look like exploits. If you then decommission that mail server and reassign the IP address to another host, you should delete the suppression conditions for that source IP address.

To view or delete a defined suppression condition:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Rules**.

The Rules page appears. By default, the page lists rules alphabetically by message.

Step 4 Locate the rule or rules where you want to view or delete suppressions. You have the following options:

- To sort the current display, click on a column heading or icon. To reverse the sort, click again.
- Construct a filter by clicking on keywords or arguments in the filter panel on the left. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy, page 27-9 and Setting a Rule Filter in an Intrusion Policy, page 27-17.

  The page refreshes to display all matching rules.

Step 5 Select the rule or rules for which you want to view or delete suppressions. You have the following options:

- To select a specific rule, select the check box next to the rule.
- To select all rules in the current list, select the check box at the top of the column.

Step 6 You have two options:

- To remove all suppression for a rule, select **Event Filtering > Remove Suppressions.** Click **OK** in the confirmation pop-up window that appears.
- To remove a specific suppression setting, highlight the rule and click **Show details**. Expand the suppression settings and click **Delete** next to the suppression settings you want to remove. Click **OK** to confirm that you want to delete your selected settings.

The page refreshes and the suppression settings are deleted.

Step 7 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See Managing Intrusion Policies, page 26-3 and Editing Intrusion Policies, page 26-4 for more information.

Adding Dynamic Rule States

License: Protection

Rate-based attacks attempt to overwhelm a network or host by sending excessive traffic toward the network or host, causing it to slow down or deny legitimate requests. You can use rate-based prevention to change the action of a rule in response to excessive rule matches for specific rules.

For more information, see the following sections:

- Understanding Dynamic Rule States, page 27-28
- Setting a Dynamic Rule State, page 27-29

Understanding Dynamic Rule States

License: Protection

You can configure your intrusion policies to include a rate-based filter that detects when too many matches for a rule occur in a given time period. You can use this feature on a device deployed inline to block rate-based attacks for a specified time, then revert to a rule state where rule matches only generate events and do not drop traffic.

Rate-based attack prevention identifies abnormal traffic patterns and attempts to minimize the impact of that traffic on legitimate requests. You can identify excessive rule matches in traffic going to a particular destination IP address or addresses or coming from a particular source IP address or addresses. You can also respond to excessive matches for a particular rule across all detected traffic.

In the intrusion policy, you can configure a rate-based filter for any intrusion or preprocessor rule. The rate-based filter contains three components:

- the rule matching rate, which you configure as a count of rule matches within a specific number of seconds
- a new action to be taken when the rate is exceeded, with three available actions: Generate Events, Drop and Generate Events, and Disable
- the duration of the action, which you configure as a timeout value

Note that when started, the new action occurs until the timeout is reached, even if the rate falls below the configured rate during that time period. When the timeout is reached, if the rate has fallen below the threshold, the action for the rule reverts to the action initially configured for the rule.

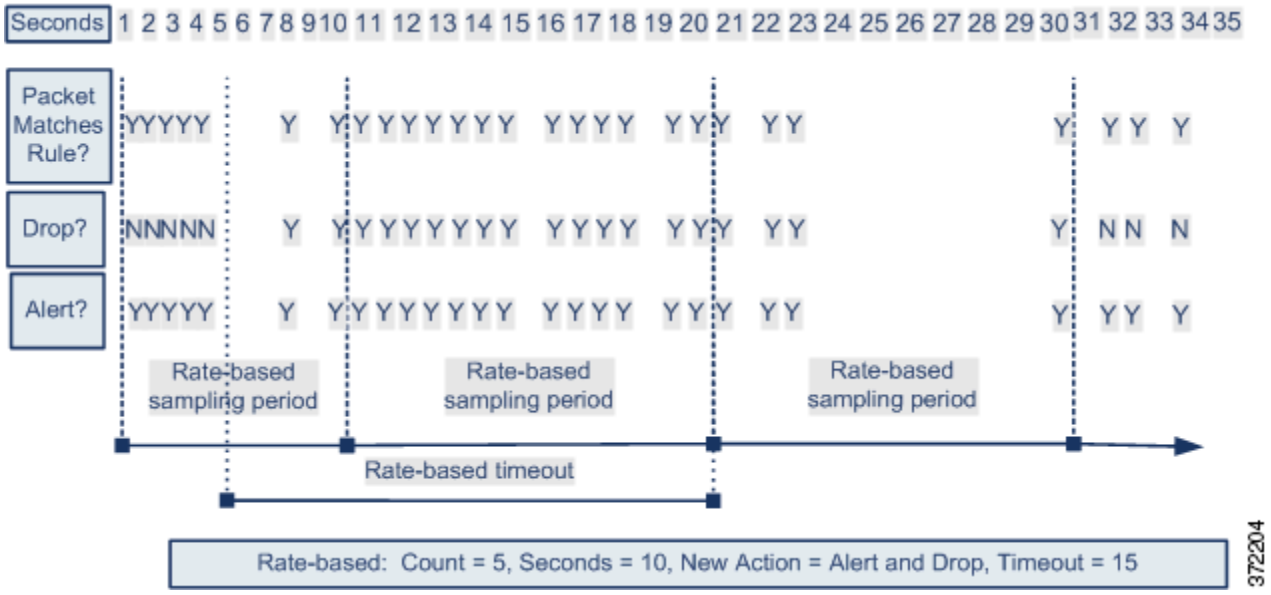
You can configure rate-based attack prevention in an inline deployment to block attacks, either temporarily or permanently. Without rate-based configuration, rules set to Generate Events do generate events, but the system does not drop packets for those rules. However, if the attack traffic matches rules that have rate-based criteria configured, the rate action may cause packet dropping to occur for the period of time that the rate action is active, even if those rules are not initially set to Drop and Generate Events.

Note Rate-based actions cannot enable disabled rules or drop traffic that matches disabled rules.

You can define multiple rate-based filters on the same rule. The first filter listed in the intrusion policy has the highest priority. Note that when two rate-based filter actions conflict, the action of the first rate-based filter is carried out.

The following diagram shows an example where an attacker is attempting to access a host. Repeated attempts to find a password trigger a rule which has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events after rule matches occur five times in a 10-second span. The new rule attribute times out after 15 seconds.

After the timeout, note that packets are still dropped in the rate-based sampling period that follows. If the sampled rate is above the threshold in the current or previous sampling period, the new action continues. The new action reverts to Generate Events only after a sampling period completes where the sampled rate was below the threshold rate.



Setting a Dynamic Rule State

License: Protection

In some cases, you may not want to set a rule to the Drop and Generate Events state because you do not want to drop every packet that matches the rule, but you do want to drop packets matching the rule if a particular rate of matches occurs in a specified time. Dynamic rule states let you configure the rate that should trigger a change in the action for a rule, what the action should change to when the rate is met, and how long the new action should persist.

You set the number of hits for that rule by specifying a count and the number of seconds within which that number of hits should occur to trigger the action change. In addition, you can set a timeout to cause the action to revert to the previous state for the rule when the timeout expires.

You can define multiple dynamic rule state filters for the same rule. The first filter listed in the rule details in the intrusion policy has the highest priority. Note that when two rate-based filter actions conflict, the action of the first rate-based filter is carried out.

Note that a revert icon ( ) appears in a field when you type an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

Note Dynamic rule states cannot enable disabled rules or drop traffic that matches disabled rules.

To add a dynamic rule state:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Rules**.

The Rules page appears.

Step 4 Locate the rule or rules where you want to add a dynamic rule state. You have the following options:

- To sort the current display, click on a column heading or icon. To reverse the sort, click again.
- Construct a filter by clicking on keywords or arguments in the filter panel on the left. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy, page 27-9 and Setting a Rule Filter in an Intrusion Policy, page 27-17.

  The page refreshes to display all matching rules.

Step 5 Select the rule or rules where you want to add a dynamic rule state. You have the following options:

- To select a specific rule, select the check box next to the rule.
- To select all the rules in the current list, select the check box at the top of the column.

Step 6 Select **Dynamic State > Add Rate-Based Rule State**.

The Add Rate-Based Rule State dialog box appears.

Step 7 From the **Track By** drop-down list, select how you want the rule matches tracked:

- Select **Source** to track the number of hits for that rule from a specific source or set of sources.
- Select **Destination** to track the number of hits for that rule to a specific destination or set of destinations.
- Select **Rule** to track all matches for that rule.

Step 8 If you set **Track By** to **Source** or **Destination**, enter the address of each host you want to track in the **Network** field.

You can specify a single IP address, address block, variable, or a comma-separated list comprised of any combination of these. For information on using IPv4 CIDR and IPv6 prefix length address blocks, see IP Address Conventions, page 1-4.

Step 9 Next to **Rate**, indicate the number of rule matches per time period to set the attack rate:

- In the **Count** field, using an integer between 1 and 2147483647, specify the number of rule matches you want to use as your threshold.
- In the **Seconds** field, using an integer between 1 and 2147483647, specify the number of seconds that make up the time period for which attacks are tracked.

Step 10 From the **New State** drop-down list, specify the new action to be taken when the conditions are met:

- Select **Generate Events** to generate an event.
- Select **Drop and Generate Events** to generate an event and drop the packet that triggered the event in inline deployments or generate an event in passive deployments.
- Select **Disabled** to take no action.

Step 11 In the **Timeout** field, type the number of seconds you want the new action to remain in effect. After the timeout occurs, the rule reverts to its original state. Specify `0` or leave the **Timeout** field blank to prevent the new action from timing out.

Step 12 Click **OK**.

The system adds the dynamic rule state and displays a dynamic state icon ( ) next to the rule in the Dynamic State column. If you add multiple dynamic rule state filters to a rule, a number over the icon indicates the number of filters.

If any required fields are blank, you receive an error message indicating which fields you must fill.

Tip To delete all dynamic rule settings for a set of rules, select the rules on the Rules page, then select **Dynamic State > Remove Rate-Based States**. You can also delete individual rate-based rule state filters from the rule details for the rule by selecting the rule, clicking **Show details**, then clicking **Delete** by the rate-based filter you want to remove.

Step 13 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache.

See Managing Intrusion Policies, page 26-3 and Editing Intrusion Policies, page 26-4 for more information.

Adding SNMP Alerts

License: Protection

If you configure SNMP alerting for your ASA FirePOWER module, you can configure specific rules to provide an SNMP alert when the rule generates an event. For more information, see Using SNMP Responses, page 39-1.

To set an SNMP alert:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Rules**.

The Rules page appears.

Step 4 Locate the rule or rules where you want to set SNMP alerts. You have the following options:

- To sort the current display, click on a column heading or icon. To reverse the sort, click again.
- Construct a filter by clicking on keywords or arguments in the filter panel on the left. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy, page 27-9 and Setting a Rule Filter in an Intrusion Policy, page 27-17.

  The page refreshes to display all matching rules.

Step 5 Select the rule or rules where you want to set SNMP alerts:

- To select a specific rule, select the check box next to the rule.
- To select all the rules in the current list, select the check box at the top of the column.

Step 6 Select **Alerting > Add SNMP Alert**.

The system adds the alert and displays an alert icon ( ) next to the rule in the Alerting column. If you add multiple alert types to a rule, a number over the icon indicates the number of alert types.

Tip To remove an SNMP alert from a rule, click the check box next to the rule and select **Alerting > Remove SNMP Alerts**, then click **OK** to confirm the deletion.

Step 7 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See Managing Intrusion Policies, page 26-3 and Editing Intrusion Policies, page 26-4 for more information.

Adding Rule Comments

License: Protection

You can add comments to a rule. Any comments you add can be seen in the Rule Details view on the Rules page.

After you commit the intrusion policy changes containing the comment, you can also view the comment by clicking **Rule Comment** on the rule Edit page. For more information on editing rules, see Modifying Existing Rules, page 30-103.

To add a comment to a rule:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Rules**.

The Rules page appears.

Step 4 Locate the rule or rules where you want to add a comment to a rule. You have the following options:

- To sort the current display, click on a column heading or icon. To reverse the sort, click again.
- Construct a filter by clicking on keywords or arguments in the filter panel on the left. For more information, see the following topics: Understanding Rule Filtering in an Intrusion Policy, page 27-9 and Setting a Rule Filter in an Intrusion Policy, page 27-17.

  The page refreshes to display all matching rules.

Step 5 Select the rule or rules where you want to add a comment:

- To select a specific rule, select the check box next to the rule.
- To select all the rules in the current list, select the check box at the top of the column.

Step 6 Select **Comments > Add Rule Comment**.

The Add Comment dialog box appears.

Step 7 In the **Comment** field, type the rule comment.

Step 8 Click **OK**.

The system adds the comment and displays a comment icon ( ) next to the rule in the Comments column. If you add multiple comments to a rule, a number over the icon indicates the number of comments.

Tip To delete a rule comment, highlight the rule and click **Show Details**, then click **Delete** in the **Comments** section. Note that you can only delete a comment if the comment is cached with uncommitted intrusion policy changes. After intrusion policy changes are committed, the rule comment is permanent.

Step 9 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache.

See Managing Intrusion Policies, page 26-3 and Editing Intrusion Policies, page 26-4 for more information.



Detecting Specific Threats

You can use several preprocessors in a network analysis policy to detect specific threats to your monitored network, such as Back Orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic. Note that when an intrusion rule or rule argument requires a disabled preprocessor, the system automatically uses it with its current configuration even though it remains disabled in the network analysis policy's user interface. For more information, see Limitations of Custom Policies, page 18-11.

You can also use sensitive data detection, which you configure in an intrusion policy, to detect unsecured transmission of sensitive numerical data.

See the following sections for more information on detecting specific threats:

- Detecting Back Orifice, page 28-1 explains detection of Back Orifice attacks.
- Detecting Portscans, page 28-3 describes the different types of portscans and explains how you can use portscan detection to identify threats to your networks before they develop into attacks.
- Preventing Rate-Based Attacks, page 28-9 explains how to limit denial of service (DoS) and SYN flood attacks.
- Detecting Sensitive Data, page 28-19 explains how to detect and generate events on sensitive data such as credit card numbers and Social Security numbers in ASCII text.

Detecting Back Orifice

License: Protection

The ASA FirePOWER module provides a preprocessor that detects the existence of the Back Orifice program. This program can be used to gain admin access to your Windows hosts. The Back Orifice preprocessor analyzes UDP traffic for the Back Orifice magic cookie, "`*!*QWTY?`", which is located in the first eight bytes of the packet and is XOR-encrypted.

The Back Orifice preprocessor has a configuration page, but no configuration options. When it is enabled, you must also enable the preprocessor rules in the following table for the preprocessor to generate corresponding events.

Table 28-1 Back Orifice GID:SIDs

Preprocessor rule GID:SID	Description
105:1	Back Orifice traffic detected
105:2	Back Orifice client traffic detected
105:3	Back Orifice server traffic detected
105:4	Back Orifice snort buffer attack detected

To view the Back Orifice Detection page:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 In the navigation panel on the left, click **Settings**.

The Settings page appears.

Step 8 You have two choices, depending on whether **Back Orifice Detection** under **Specific Threat Detection** is enabled:

- If the preprocessor is enabled, click **Edit**.
- If the preprocessor is disabled, click **Enabled**, then click **Edit**.

The Back Orifice Detection page appears. A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Detecting Portscans

License: Protection

A portscan is a form of network reconnaissance that is often used by attackers as a prelude to an attack. In a portscan, an attacker sends specially crafted packets to a targeted host. By examining the packets that the host responds with, the attacker can often determine which ports are open on the host and, either directly or by inference, which application protocols are running on these ports.

Note that when portscan detection is enabled, you must enable rules on the intrusion policy Rules page with generator ID (GID) 122 for enabled portscan types for the portscan detector to generate portscan events. See Setting Rule States, page 27-19 and Table 28-5 on page 28-7 for more information.

By itself, a portscan is not evidence of an attack. In fact, some of the portscanning techniques used by attackers can also be employed by legitimate users on your network. Cisco's portscan detector is designed to help you determine which portscans might be malicious by detecting patterns of activity.

Attackers are likely to use several methods to probe your network. Often they use different protocols to draw out different responses from a target host, hoping that if one type of protocol is blocked, another may be available. The following table describes the protocols you can activate in the portscan detector.

Table 28-2 Protocol Types

Protocol	Description
TCP	Detects TCP probes such as SYN scans, ACK scans, TCP connect() scans, and scans with unusual flag combinations such as Xmas tree, FIN, and NULL
UDP	Detects UDP probes such as zero-byte UDP packets
ICMP	Detects ICMP echo requests (pings)
IP	Detects IP protocol scans. These scans differ from TCP and UDP scans because the attacker, instead of looking for open ports, is trying to discover which IP protocols are supported on a target host.

Note For events generated by the portscan connection detector, the protocol number is set to 255. Because portscan does not have a specific protocol associated with it by default, the Internet Assigned Numbers Authority (IANA) does not have a protocol number assigned to it. IANA designates 255 as a reserved number, so that number is used in portscan events to indicate that there is not an associated protocol for the event.

Portscans are generally divided into four types based on the number of targeted hosts, the number of scanning hosts, and the number of ports that are scanned. The following table describes the kinds of portscan activity you can detect.

Table 28-3 Portscan Types

Type	Description
Portscan Detection	A one-to-one portscan in which an attacker uses one or a few hosts to scan multiple ports on a single target host. One-to-one portscans are characterized by: • a low number of scanning hosts • a single host that is scanned • a high number of ports scanned This option detects TCP, UDP, and IP portscans.
Port Sweep	A one-to-many portsweep in which an attacker uses one or a few hosts to scan a single port on multiple target hosts. Portsweeps are characterized by: • a low number of scanning hosts • a high number of scanned hosts • a low number of unique ports scanned This option detects TCP, UDP, ICMP, and IP portsweeps.
Decoy Portscan	A one-to-one portscan in which the attacker mixes spoofed source IP addresses with the actual scanning IP address. Decoy portscans are characterized by: • a high number of scanning hosts • a low number of ports that are scanned only once • a single (or a low number of) scanned hosts The decoy portscan option detects TCP, UDP, and IP protocol portscans.
Distributed Portscan	A many-to-one portscan in which multiple hosts query a single host for open ports. Distributed portscans are characterized by: • a high number of scanning hosts • a high number of ports that are scanned only once • a single (or a low number of) scanned hosts The distributed portscan option detects TCP, UDP, and IP protocol portscans.

The information that the portscan detector learns about a probe is largely based on seeing negative responses from the probed hosts. For example, when a web client tries to connect to a web server, the client uses port 80/tcp and the server can be counted on to have that port open. However, when an attacker probes a server, the attacker does not know in advance if it offers web services. When the portscan detector sees a negative response (that is, an ICMP unreachable or TCP RST packet), it records the response as a potential portscan. The process is more difficult when the targeted host is on the other side of a device such as a firewall or router that filters negative responses. In this case, the portscan detector can generate *filtered* portscan events based on the sensitivity level that you select.

The following table describes the three different sensitivity levels you can choose from.

Table 28-4 Sensitivity Levels

Level	Description
Low	Detects only negative responses from targeted hosts. Select this sensitivity level to suppress false positives, but keep in mind that some types of portscans (slow scans, filtered scans) might be missed. This level uses the shortest time window for portscan detection.
Medium	Detects portscans based on the number of connections to a host, which means that you can detect filtered portscans. However, very active hosts such as network address translators and proxies may generate false positives. Note that you can add the IP addresses of these active hosts to the **Ignore Scanned** field to mitigate this type of false positive. This level uses a longer time window for portscan detection.
High	Detects portscans based on a time window, which means that you can detect time-based portscans. However, if you use this option, you should be careful to tune the detector over time by specifying IP addresses in the **Ignore Scanned** and **Ignore Scanner** fields. This level uses a much longer time window for portscan detection.

See the following sections for more information:

- Configuring Portscan Detection, page 28-5
- Understanding Portscan Events, page 28-7

Configuring Portscan Detection

License: Protection

The portscan detection configuration options allow you to finely tune how the portscan detector reports scan activity.

Note that when portscan detection is enabled, you must enable rules on the Rules page with generator ID (GID) 122 for enabled portscan types for the portscan detector to generate portscan events. See Setting Rule States, page 27-19 and the Portscan Detection SIDs (GID:122) table for more information.

To configure portscan detection:

Admin/Intrusion Admin

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See [was Committing Intrusion Policy Changes; update xref] for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 In the navigation panel on the left, click **Settings**.

The Settings page appears.

Step 8 You have two choices, depending on whether **Portscan Detection** under **Specific Threat Detection** is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

  The Portscan Detection page appears. A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 In the **Protocol** field, specify which of the following protocols you want to enable:

- TCP
- UDP
- ICMP
- IP

Use Ctrl or Shift while clicking to select multiple protocols or clear individual protocols. See the Protocol Types table for more information.

Note that you must ensure TCP stream processing is enabled to detect scans over TCP, and that UDP stream processing is enabled to detect scans over UDP.

Step 10 In the **Scan Type** field, specify which of the following portscans you want to detect:

- Portscan Detection
- Port Sweep
- Decoy Portscan
- Distributed Portscan

Use Ctrl or Shift while clicking to select or deselect multiple protocols. See the Portscan Types table for more information.

Step 11 In the **Sensitivity Level** list, select the level you want to use: low, medium, or high.

See the Sensitivity Levels table for more information.

Step 12 Optionally, in the **Watch IP** field, specify which host you want to watch for signs of portscan activity, or leave the field blank to watch all network traffic.

You can specify a single IP address or address block, or a comma-separated lists of either or both. For information on using IPv4 and IPv6 address blocks, see IP Address Conventions, page 1-4.

Step 13 Optionally, in the **Ignore Scanners** field, specify which hosts you want to ignore as scanners. Use this field to indicate hosts on your network that are especially active. You may need to modify this list of hosts over time.

You can specify a single IP address or address block, or a comma-separated lists of either or both. For information on using IPv4 and IPv6 address blocks, see IP Address Conventions, page 1-4.

Step 14 Optionally, in the **Ignore Scanned** field, specify which hosts you want to ignore as the target of a scan. Use this field to indicate hosts on your network that are especially active. You may need to modify this list of hosts over time.

You can specify a single IP address or address block, or a comma-separated lists of either or both. For information on using IPv4 and IPv6 address blocks, see IP Address Conventions, page 1-4.

Step 15 Optionally, clear the **Detect Ack Scans** check box to discontinue monitoring of sessions picked up in mid-stream.

Note Detection of mid-stream sessions helps to identify ACK scans, but may cause false events, particularly on networks with heavy traffic and dropped packets.

Step 16 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Understanding Portscan Events

License: Protection

When portscan detection is enabled, you must enable rules with generator ID (GID) 122 and a Snort® ID (SID) from among SIDs 1 through 27 to generate events for each enabled portscan type. See Setting Rule States, page 27-19 for more information. The **Preprocessor Rule SID** column in the following table lists the SID for the preprocessor rule you must enable for each portscan type.

Table 28-5 Portscan Detection SIDs (GID:122)

Portscan Type	Protocol:	Sensitivity Level	Preprocessor Rule SID
Portscan Detection	TCP	Low	1
		Medium or High	5
	UDP	Low	17
		Medium or High	21
	ICMP	Low	Does not generate events.
		Medium or High	Does not generate events.
	IP	Low	9
		Medium or High	13
Port Sweep	TCP	Low	3, 27
		Medium or High	7
	UDP	Low	19
		Medium or High	23
	ICMP	Low	25
		Medium or High	26
	IP	Low	11
		Medium or High	15

Table 28-5 Portscan Detection SIDs (GID:122) (continued)

Portscan Type	Protocol:	Sensitivity Level	Preprocessor Rule SID
Decoy Portscan	TCP	Low	2
		Medium or High	6
	UDP	Low	18
		Medium or High	22
	ICMP	Low	Does not generate events.
		Medium or High	Does not generate events.
	IP	Low	10
		Medium or High	14
Distributed Portscan	TCP	Low	4
		Medium or High	8
	UDP	Low	20
		Medium or High	24
	ICMP	Low	Does not generate events.
		Medium or High	Does not generate events.
	IP	Low	12
		Medium or High	16

When you enable the accompanying preprocessor rules, the portscan detector generates intrusion events that you can view just as you would any other intrusion event. However, the information presented on the packet view is different from the other types of intrusion events. This section describes the fields that appear on the packet view for a portscan event and how you can use that information to understand the types of probes that occur on your network.

Begin by using the intrusion event views to drill down to the packet view for a portscan event.

Note that you cannot download a portscan packet because single portscan events are based on multiple packets; however, the portscan packet view provides all usable packet information.



Note For events generated by the portscan connection detector, the protocol number is set to 255. Because portscan does not have a specific protocol associated with it by default, the Internet Assigned Numbers Authority (IANA) does not have a protocol number assigned to it. IANA designates 255 as a reserved number, so that number is used in portscan events to indicate that there is not an associated protocol for the event.

The following table describes the information provided in the packet view for portscan events.

Table 28-6 Portscan Packet View

Information	Description
Device	The device that detected the event.
Time	The time when the event occurred.
Message	The event message generated by the preprocessor.
Source IP	The IP address of the scanning host.
Destination IP	The IP address of the scanned host.

Table 28-6 Portscan Packet View (continued)

Information	Description
Priority Count	The number of negative responses (for example, TCP RSTs and ICMP unreachables) from the scanned host. The higher the number of negative responses, the higher the priority count.
Connection Count	The number of active connections on the hosts. This value is more accurate for connection-based scans such as TCP and IP.
IP Count	The number of times that the IP addresses that contact the scanned host changes. For example, if the first IP address is 10.1.1.1, the second IP is 10.1.1.2, and the third IP is 10.1.1.1, then the IP count is 3. This number is less accurate for active hosts such as proxies and DNS servers.
Scanner/Scanned IP Range	The range of IP addresses for the scanned hosts or the scanning hosts, depending on the type of scan. For portsweeps, this field shows the IP range of scanned hosts. For portscans, this shows the IP range of the scanning hosts.
Port/Proto Count	For TCP and UDP portscans, the number of times that the port being scanned changes. For example, if the first port scanned is 80, the second port scanned is 8080, and the third port scanned is again 80, then the port count is 3. For IP protocol portscans, the number of times that the protocol being used to connect to the scanned host changes.
Port/Proto Range	For TCP and UDP portscans, the range of the ports that were scanned. For IP protocol portscans, the range of IP protocol numbers that were used to attempt to connect to the scanned host.
Open Ports	The TCP ports that were open on the scanned host. This field appears only when the portscan detects one or more open ports.

Preventing Rate-Based Attacks

License: Protection

Rate-based attacks are attacks that depend on frequency of connection or repeated attempts to perpetrate the attack. You can use rate-based detection criteria to detect a rate-based attack as it occurs and respond to it when it happens, then return to normal detection settings after it stops. For more information on configuring rate-based detection, see the following topics:

- Understanding Rate-Based Attack Prevention, page 28-9
- Rate-Based Attack Prevention and Other Filters, page 28-12
- Configuring Rate-Based Attack Prevention, page 28-17
- Understanding Dynamic Rule States, page 27-28
- Setting a Dynamic Rule State, page 27-29

Understanding Rate-Based Attack Prevention

License: Protection

You can configure your network analysis policy to include rate-based filters that detect excessive activity directed at hosts on your network. You can use this feature on a device deployed in inline mode to block rate-based attacks for a specified time, then revert to only generating events and not drop traffic.

Rate-based attack prevention identifies abnormal traffic patterns and attempts to minimize the impact of that traffic on legitimate requests. Rate-based attacks usually have one of the following characteristics:

- any traffic containing excessive incomplete connections to hosts on the network, indicating a SYN flood attack

  To configure SYN attack detection, see Preventing SYN Attacks, page 28-11.

- any traffic containing excessive complete connections to hosts on the network, indicating a TCP/IP connection flood attack

  To configure simultaneous connection detection, see Controlling Simultaneous Connections, page 28-12.

- excessive rule matches in traffic going to a particular destination IP address or addresses or coming from a particular source IP address or addresses.

  To configure source or destination-based dynamic rule states, see Setting a Dynamic Rule State, page 27-29.

- excessive matches for a particular rule across all traffic.

  To configure rule-based dynamic rule states, see Setting a Dynamic Rule State, page 27-29.

In a network analysis policy, you can either configure SYN flood or TCP/IP connection flood detection for the entire policy; in an intrusion policy, you can set rate-based filters for individual intrusion or preprocessor rules. Note that manually adding a rate-based filter to rules 135:1 and 135:2 has no effect. Rules with GID:135 use the client as the source value and the server as the destination value. See Preventing SYN Attacks, page 28-11 and Controlling Simultaneous Connections, page 28-12 for more information.

Each rate-based filter contains several components:

- for policy-wide or rule-based source or destination settings, the network address designation
- the rule matching rate, which you configure as a count of rule matches within a specific number of seconds
- a new action to be taken when the rate is exceeded

  When you set a rate-based setting for the entire policy, the system generates events when it detects a rate-based attack, and optionally can drop the traffic in an inline deployment. When setting rate-based actions for individual rules, you have three available actions: Generate Events, Drop and Generate Events, and Disable.

- the duration of the action, which you configure as a timeout value

Note that when started, the new action occurs until the timeout is reached, even if the rate falls below the configured rate during that time period. When the timeout period expires, if the rate has fallen below the threshold, the action for the rule reverts to the action initially configured for the rule. For policy-wide settings, the action reverts to the action of each rule the traffic matches or stops if it does not match any rules.

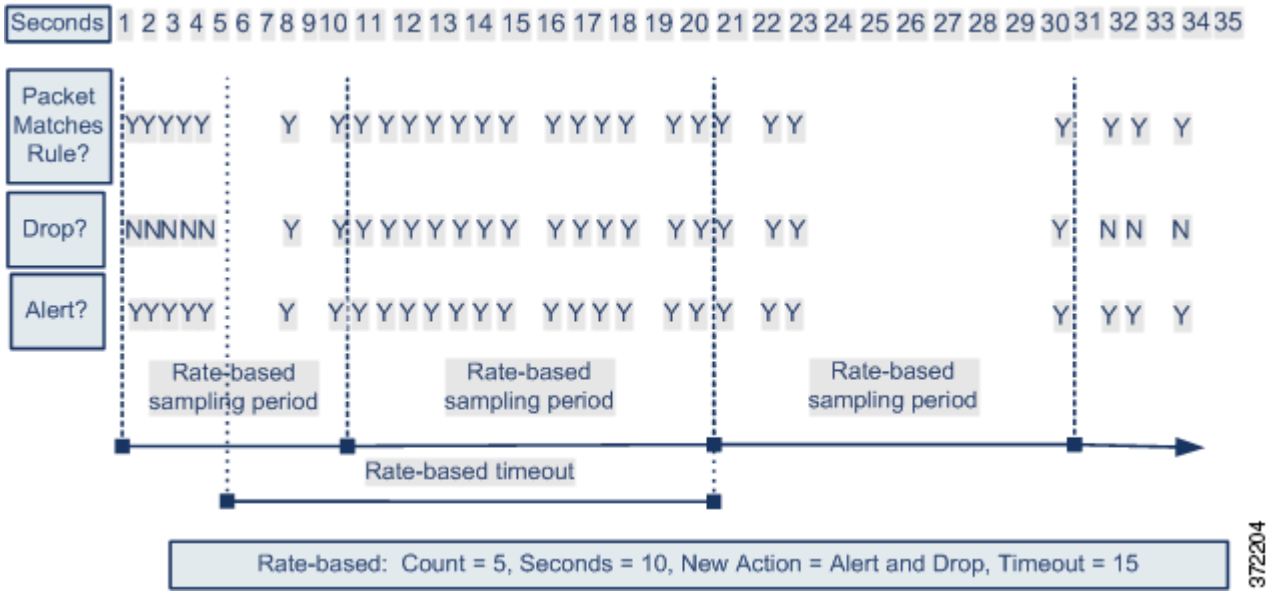
You can configure rate-based attack prevention in an inline deployment to block attacks, either temporarily or permanently. Without rate-based configuration, rules set to Generate Events create events, but the system does not drop packets for those rules. However, if the attack traffic matches rules that have rate-based criteria configured, the rate action may cause packet dropping to occur for the period of time that the rate action is active, even if those rules are not initially set to Drop and Generate Events.

Note Rate-based actions cannot enable disabled rules or drop traffic that matches disabled rules. However, if you set a rate-based filter at the policy level, you can generate events on or generate events on and drop traffic that contains an excessive number of SYN packets or SYN/ACK interactions within a designated time period.

You can define multiple rate-based filters on the same rule. The first filter listed in the intrusion policy has the highest priority. Note that when two rate-based filter actions conflict, the system implements the action of the first rate-based filter. Similarly, policy-wide rate-based filters override rate-based filters set on individual rules if the filters conflict.

The following diagram shows an example where an attacker is attempting to access a host. Repeated attempts to find a password trigger a rule which has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events after rule matches occur five times in a 10-second span. The new rule attribute times out after 15 seconds.

After the timeout, note that packets are still dropped in the rate-based sampling period that follows. If the sampled rate is above the threshold in the current or previous sampling period, the new action continues. The new action reverts to generating events only after a sampling period completes where the sampled rate is below the threshold rate.



Preventing SYN Attacks

License: Protection

The SYN attack prevention option helps you protect your network hosts against SYN floods. You can protect individual hosts or whole networks based on the number of packets seen over a period of time. If your device is deployed passively, you can generate events. If your device is placed inline, you can also drop the malicious packets. After the timeout period elapses, if the rate condition has stopped, the event generation and packet dropping stops.

For example, you could configure a setting to allow a maximum of 10 SYN packets from any one IP address, and block further connections from that IP address for 60 seconds.

Enabling this option also activates rule 135:1. Manually activating this rule has no effect. The rule state is always displayed as Disabled, and never changes. The rule generates events when this option is enabled and a defined rate condition is exceeded.

Controlling Simultaneous Connections

License: Protection

You can also limit TCP/IP connections to or from hosts on your network to prevent denial of service (DoS) attacks or excessive activity by users. When the system detects the configured number of successful connections to or from a specified IP address or range of addresses, it generates events on additional connections. The rate-based event generation continues until the timeout period elapses without the rate condition occurring. In an inline deployment you can choose to drop packets until the rate condition times out.

For example, you could configure a setting to allow a maximum of 10 successful simultaneous connections from any one IP address, and block further connections from that IP address for 60 seconds.

Enabling this option also activates rule 135:2. Manually activating this rule has no effect. The rule state is always displayed as Disabled, and never changes. The rule generates events when this option is enabled and a defined rate condition is exceeded.

Rate-Based Attack Prevention and Other Filters

License: Protection

The `detection_filter` keyword and the thresholding and suppression features provide other ways to filter either the traffic itself or the events that the system generates. You can use rate-based attack prevention alone or in any combination with thresholding, suppression, or the `detection_filter` keyword.

See the following examples for more information:

- Rate-Based Attack Prevention and Detection Filtering, page 28-12
- Dynamic Rule States and Thresholding or Suppression, page 28-13
- Policy-Wide Rate-Based Detection and Thresholding or Suppression, page 28-15
- Rate-Based Detection with Multiple Filtering Methods, page 28-16

Rate-Based Attack Prevention and Detection Filtering

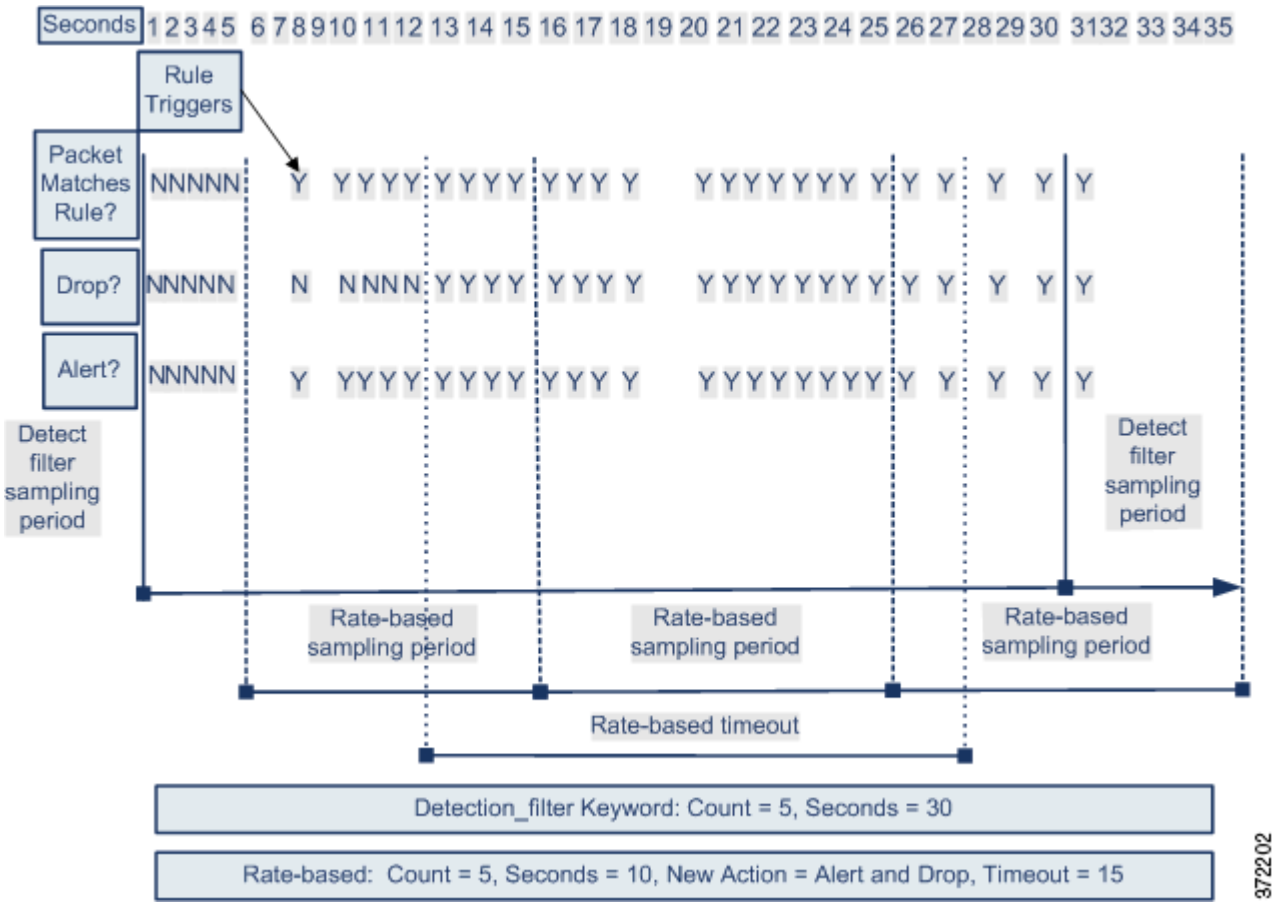
License: Protection

The `detection_filter` keyword prevents a rule from triggering until a threshold number of rule matches occur within a specified time. When a rule includes the `detection_filter` keyword, the system tracks the number of incoming packets matching the pattern in the rule per timeout period. The system can count hits for that rule from particular source or destination IP addresses. After the rate exceeds the rate in the rule, event notification for that rule begins.

The following example shows an attacker attempting a brute-force login. Repeated attempts to find a password trigger a rule that also includes the `detection_filter` keyword, with a count set to 5. This rule has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events for 20 seconds when there are five hits on the rule in a 10-second span.

As shown in the diagram, the first five packets matching the rule do not generate events because the rule does not trigger until the rate exceeds the rate indicated by the `detection_filter` keyword. After the rule triggers, event notification begins, but the rate-based criteria do not trigger the new action of Drop and Generate Events until five more packets pass.

After the rate-based criteria are met, events are generated and the packets are dropped until the rate-based timeout period expires and the rate falls below the threshold. After twenty seconds elapse, the rate-based action times out. After the timeout, note that packets are still dropped in the rate-based sampling period that follows. Because the sampled rate is above the threshold rate in the previous sampling period when the timeout happens, the rate-based action continues.



Note that although the example does not depict this, you can use the Drop and Generate Events rule state in combination with the `detection_filter` keyword to start dropping traffic when hits for the rule reach the specified rate. When deciding whether to configure rate-based settings for a rule, consider whether setting the rule to Drop and Generate Events and including the `detection_filter` keyword would achieve the same result, or whether you want to manage the rate and timeout settings in the intrusion policy. For more information, see Setting Rule States, page 27-19.

Dynamic Rule States and Thresholding or Suppression

License: Protection

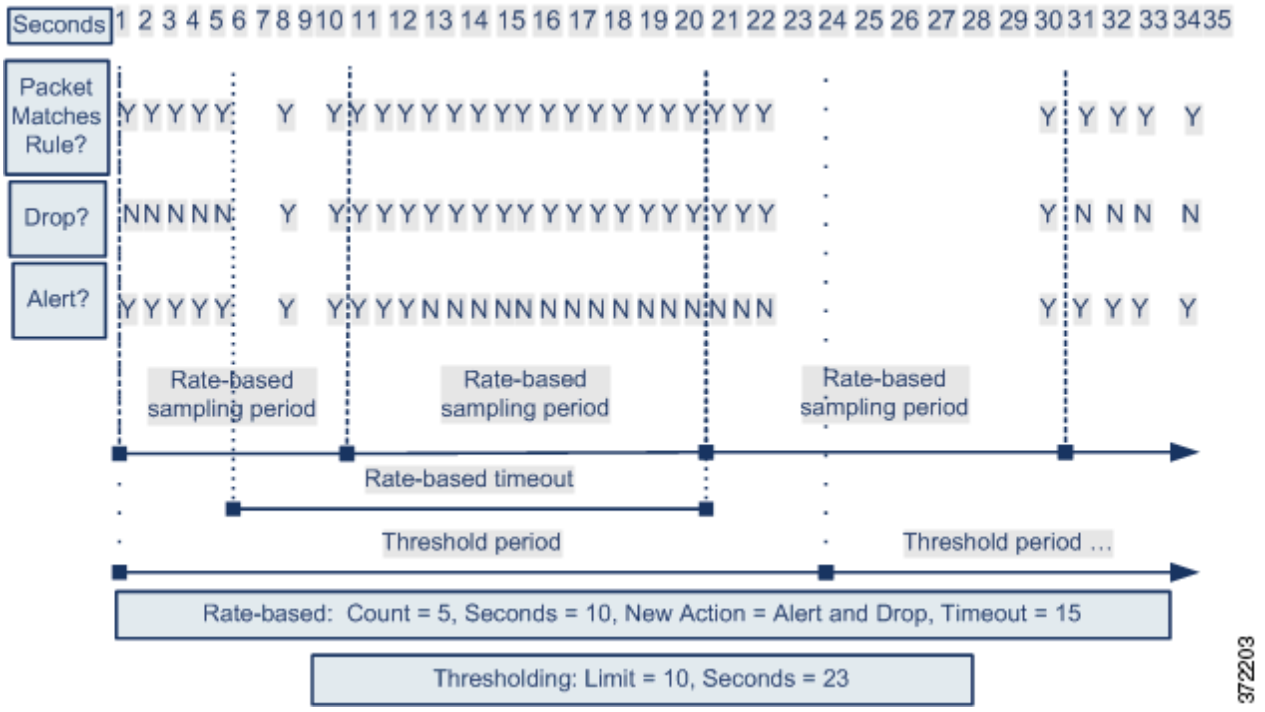
You can use thresholding and suppression to reduce excessive events by limiting the number of event notifications for a rule or by suppressing notifications altogether for that rule. For more information on the available options for thresholding and suppression, see Configuring Event Thresholding, page 27-21 and Configuring Suppression Per Intrusion Policy, page 27-25.

If you apply suppression to a rule, the system suppresses event notifications for that rule for all applicable IP addresses even if a rate-based action change occurs. However, the interaction between thresholding and rate-based criteria is more complex.

The following example shows an attacker attempting a brute-force login. Repeated attempts to find a password trigger a rule that has rate-based attack prevention configured. The rate-based settings change the rule attribute to Drop and Generate Events for 15 seconds when there are five hits on the rule in 10 seconds. In addition, a limit threshold limits the number of events the rule can generate to 10 events in 23 seconds.

As shown in the diagram, the rule generates events for the first five matching packets. After five packets, the rate-based criteria trigger the new action of Drop and Generate Events, and for the next five packets the rule generates events and the system drops the packet. After the tenth packet, the limit threshold has been reached, so for the remaining packets the system does not generate events but does drop the packets.

After the timeout, note that packets are still dropped in the rate-based sampling period that follows. If the sampled rate is above the threshold rate in the current or previous sampling period, the new action continues. The new action reverts to Generate Events only after a sampling period completes where the sampled rate is below the threshold rate.



Note that although it is not shown in this example, if a new action triggers because of rate-based criteria *after* a threshold has been reached, the system generates a single event to indicate the change in action. So, for example, when the limit threshold of 10 is reached and the system stops generating events and the action changes from Generate Events to Drop and Generate Events on the 14th packet, the system generates an eleventh event to indicate the change in action.

Policy-Wide Rate-Based Detection and Thresholding or Suppression

License: Protection

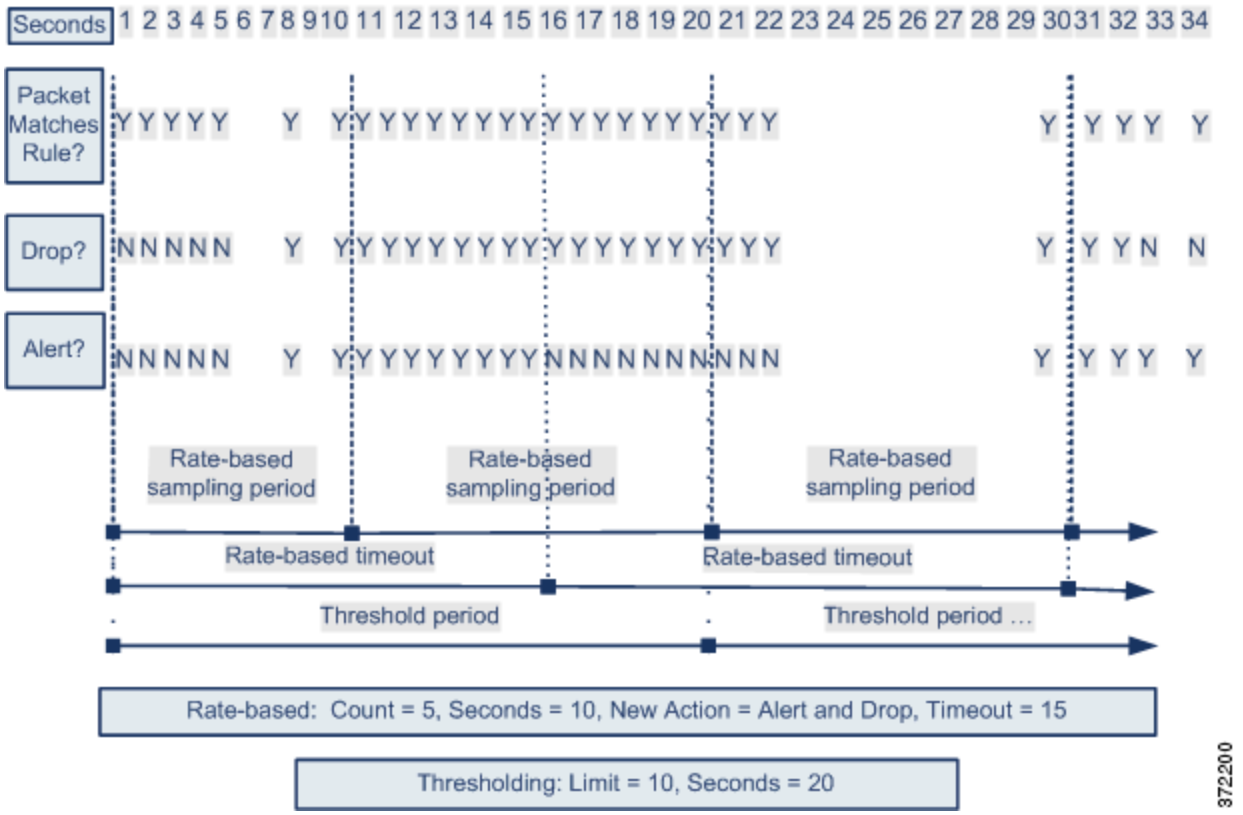
You can use thresholding and suppression to reduce excessive events by limiting the number of event notifications for a source or destination or by suppressing notifications altogether for that rule. For more information on the available options for thresholding and suppression, see Configuring Global Thresholds, page 29-3, Configuring Event Thresholding, page 27-21, and Configuring Suppression Per Intrusion Policy, page 27-25.

If suppression is applied to a rule, event notifications for that rule for all applicable IP addresses are suppressed even if a rate-based action change occurs because of a policy-wide or rule-specific rate-based setting. However, the interaction between thresholding and rate-based criteria is more complex.

The following example shows an attacker attempting denial of service (DoS) attacks on hosts in your network. Many simultaneous connections to hosts from the same sources trigger a policy-wide Control Simultaneous Connections setting. The setting generates events and drops malicious traffic when there are five connections from one source in 10 seconds. In addition, a global limit threshold limits the number of events any rule or setting can generate to 10 events in 20 seconds.

As shown in the diagram, the policy-wide setting generates events for the first ten matching packets and drops the traffic. After the tenth packet, the limit threshold is reached, so for the remaining packets no events are generated but the packets are dropped.

After the timeout, note that packets are still dropped in the rate-based sampling period that follows. If the sampled rate is above the threshold rate in the current or previous sampling period, the rate-based action of generating events and dropping traffic continues. The rate-based action stops only after a sampling period completes where the sampled rate is below the threshold rate.



Note that although it is not shown in this example, if a new action triggers because of rate-based criteria *after* a threshold has been reached, the system generates a single event to indicate the change in action. So, for example, if the limit threshold of 10 has been reached and the system stops generating events and the action changes to Drop and Generate events on the 14th packet, the system generates an eleventh event to indicate the change in action.

Rate-Based Detection with Multiple Filtering Methods

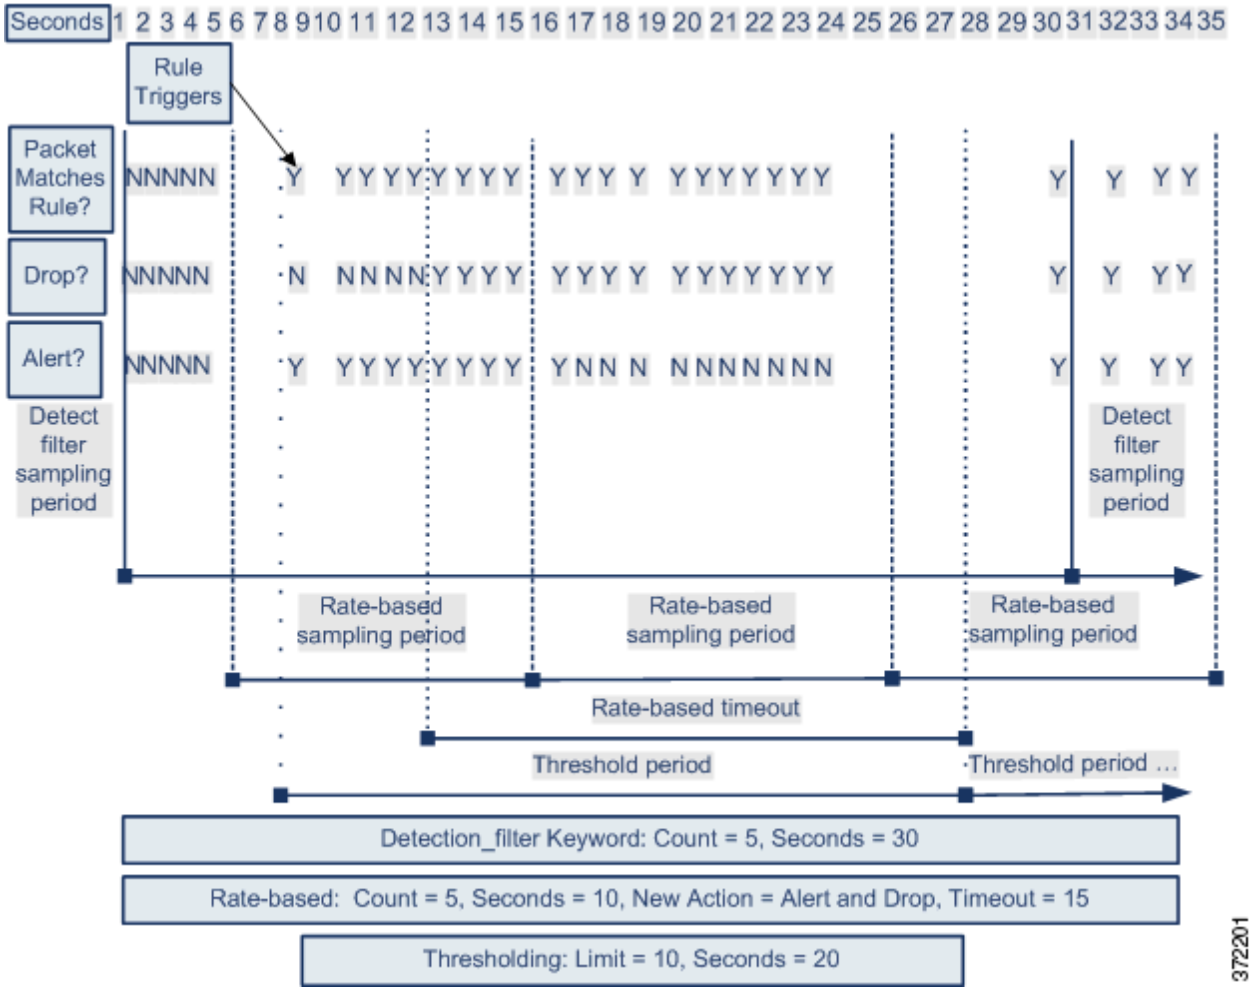
License: Protection

You may encounter situations where the `detection_filter` keyword, thresholding or suppression, and rate-based criteria all apply to the same traffic. When you enable suppression for a rule, events are suppressed for the specified IP addresses even if a rate-based change occurs.

The following example shows an attacker attempting a brute force login, and describes a case where a `detection_filter` keyword, rate-based filtering, and thresholding interact. Repeated attempts to find a password trigger a rule which includes the `detection_filter` keyword, with a count set to 5. This rule also has rate-based attack prevention settings that change the rule attribute to Drop and Generate Events for 30 seconds when there are five rule hits in 15 seconds. In addition, a limit threshold limits the rule to 10 events in 30 seconds.

As shown in the diagram, the first five packets matching the rule do not cause event notification because the rule does not trigger until the rate indicated in the `detection_filter` keyword is exceeded. After the rule triggers, event notification begins, but the rate-based criteria do not trigger the new action of Drop and Generate Events until five more packets pass. After the rate-based criteria are met, the system generates events for packets 11-15 and drops the packets. After the fifteenth packet, the limit threshold has been reached, so for the remaining packets the system does not generate events but does drop the packets.

After the rate-based timeout, note that packets are still dropped in the rate-based sampling period that follows. Because the sampled rate is above the threshold rate in the previous sampling period, the new action continues.



Configuring Rate-Based Attack Prevention

License: Protection

You can configure rate-based attack prevention at the policy level to stop SYN flood attacks. You can also stop excessive connections from a specific source or to a specific destination.

To configure rate-based attack prevention:

Admin/Intrusion Admin

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to edit.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

The access control policy advanced settings page appears.

Step 4 Click the edit icon ( ) next to **Network Analysis and Intrusion Policies**.

The Network Analysis and Intrusion Policies pop-up window appears.

Step 5 Click **Network Analysis Policy List**.

The Network Analysis Policy List pop-up window appears.

Step 6 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 7 In the navigation panel on the left, click **Settings**.

The Settings page appears.

Step 8 You have two choices, depending on whether **Rate-Based Attack Prevention** under **Specific Threat Detection** is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Rate-Based Attack Prevention page appears. A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 9 You have two options:

- To prevent incomplete connections intended to flood a host, click **Add** under **SYN Attack Prevention**.

  The SYN Attack Prevention dialog box appears.

- To prevent excessive numbers of connections, click **Add** under **Control Simultaneous Connections**.

  The Control Simultaneous Connections dialog box appears.

Step 10 Select how you want to track traffic:

- To track all traffic from a specific source or range of sources, select **Source** from the **Track By** drop-down list and type a single IP address or address block in the **Network** field.
- To track all traffic to a specific destination or range of destinations, select **Destination** from the **Track By** drop-down list and type an IP address or address block in the **Network** field.

Note that the system tracks traffic separately for each IP address included in the Network field. Traffic from an IP address that exceeds the configured rate results in generated events only for that IP address. As an example, you might set a source CIDR block of `10.1.0.0/16` for the network setting and configure the system to generate events when there are ten simultaneous connections open. If eight connections are open from 10.1.4.21 and six from 10.1.5.10, the system does not generate events, because neither source has the triggering number of connections open. However, if eleven simultaneous connections are open from 10.1.4.21, the system generates events only for the connections from 10.1.4.21.

For information on using CIDR notation and prefix lengths, see IP Address Conventions, page 1-4.

Step 11 Indicate the triggering rate for the rate tracking setting:

- For SYN attack configuration, indicate the number of SYN packets per number of seconds in the **Rate** fields.
- For simultaneous connection configuration, indicate the number of connections in the **Count** field.

Step 12 To drop packets matching the rate-based attack prevention settings, select **Drop**.

Step 13 In the **Timeout** field, indicate the time period after which to stop generating events, and if applicable, dropping, for traffic with the matching pattern of SYNs or simultaneous connections.

Caution Timeout values can be integers from 1 to 1,000,000. However, setting a high timeout value may entirely block connection to a host in an inline deployment.

Step 14 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Detecting Sensitive Data

License: Protection

Sensitive data such as Social Security numbers, credit card numbers, driver's license numbers, and so on may be leaked onto the Internet, intentionally or accidentally. The system provides a sensitive data preprocessor that can detect and generate events on sensitive data in ASCII text, which can be particularly useful in detecting accidental data leaks.

The system does not detect encrypted or obfuscated sensitive data, or sensitive data in a compressed or encoded format such as a Base64-encoded email attachment. For example, the system would detect the phone number (555)123-4567, but not an obfuscated version where each number is separated by spaces, as in (5 5 5) 1 2 3 - 4 5 6 7, or by intervening HTML code, such as (555)-123-4567. However, the system would detect, for example, the HTML coded number (555)-123-4567 where no intervening codes interrupt the numbering pattern.

Tip The sensitive data preprocessor can detect sensitive data in unencrypted Microsoft Word files that are uploaded and downloaded using FTP or HTTP; this is possible because of the way Word files group ASCII text and formatting commands separately.

The system detects sensitive data per TCP session by matching individual data types against traffic. You can modify the default settings for each data type and for global options that apply to all data types in your intrusion policy. Cisco provides predefined, commonly used data types. You can also create custom data types.

A sensitive data preprocessor rule is associated with each data type. You enable sensitive data detection and event generation for each data type by enabling the corresponding preprocessor rule for the data type. A link on the configuration page takes you to a filtered view of sensitive data rules on the Rules page, where you can enable and disable rules and configure other rule attributes.

When you save changes to your intrusion policy, you are given the option to automatically enable the sensitive data preprocessor if the rule associated with a data type is enabled and sensitive data detection is disabled.

See the following sections for more information:

- Deploying Sensitive Data Detection, page 28-20
- Selecting Global Sensitive Data Detection Options, page 28-20
- Selecting Individual Data Type Options, page 28-21
- Using Predefined Data Types, page 28-22

- Configuring Sensitive Data Detection, page 28-23
- Selecting Application Protocols to Monitor, page 28-25
- Special Case: Detecting Sensitive Data in FTP Traffic, page 28-26
- Using Custom Data Types, page 28-27

Deploying Sensitive Data Detection

License: Protection

Because sensitive data detection can have a high impact on the performance of your system, Cisco recommends that you adhere to the following guidelines:

- Select the No Rules Active default policy as your base intrusion policy; see Understanding System-Provided Base Policies, page 19-3 for more information.
- Ensure that the following settings are enabled in the corresponding network analysis policy:
  - **FTP and Telnet Configuration** under **Application Layer Preprocessors**
  - **IP Defragmentation** and **TCP Stream Configuration** under **Transport/Network Layer Preprocessors**.
- Apply the access control policy that includes the intrusion policy containing your sensitive data configuration to a device reserved for sensitive data detection; see Deploying Configuration Changes, page 4-11 for more information.

Selecting Global Sensitive Data Detection Options

License: Protection

Global sensitive data preprocessor options control how the preprocessor functions. You can modify global options that specify the following:

- whether the preprocessor replaces all but the last four credit card or Social Security numbers in triggering packets
- which destination hosts on your network to monitor for sensitive data
- how many total occurrences of all data types in a single session result in an event

Note that global sensitive data options are policy-specific and apply to all data types.

You can configure the following global sensitive data detection options.

Mask

Replaces with Xs all but the last four digits of credit card numbers and Social Security numbers in the triggering packet. The masked numbers appear in the intrusion event packet view in the user interface and in downloaded packets.

Networks

Specifies the destination host or hosts to monitor for sensitive data. You can specify a single IP address, address block, or a comma-separated list of either or both. The system interprets a blank field as `any`, meaning any destination IP address. For information on using IPv4 and IPv6 address blocks, see IP Address Conventions, page 1-4.

Global Threshold

Specifies the total number of all occurrences of all data types during a single session that the preprocessor must detect in any combination before generating a global threshold event. You can specify 1 through 65535.

Cisco recommends that you set the value for this option higher than the highest threshold value for any individual data type that you enable in your policy. See Selecting Individual Data Type Options, page 28-21 for more information.

Note the following points regarding global thresholds:

– You must enable preprocessor rule 139:1 to detect and generate events on combined data type occurrences. See Setting Rule States, page 27-19 for information on enabling rules in your intrusion policy.

– The preprocessor generates up to one global threshold event per session.

– Global threshold events are independent of individual data type events; that is, the preprocessor generates an event when the global threshold is reached, regardless of whether the event threshold for any individual data type has been reached, and vice versa.

Selecting Individual Data Type Options

License: Protection

Individual data types identify the sensitive data you can detect and generate events on in your specified destination network traffic. You can modify default settings for data type options that specify the following:

- a threshold that must be met for a detected data type to generate a single per-session event
- the destination ports to monitor for each data type
- the application protocols to monitor for each data type

At a minimum, each data type must specify an event threshold and at least one port or application protocol to monitor.

Each predefined data type provided by Cisco uses an otherwise inaccessible `sd_pattern` keyword to define a built-in data pattern to detect in traffic. See Table 28-8 on page 28-23 for a listing of predefined data types. You can also create custom data types for which you use simple regular expressions to specify your own data patterns. See Using Custom Data Types, page 28-27 for more information.

Note that data type names and patterns are system-wide; all other data type options are policy-specific.

The following table describes the data type options you can configure.

Table 28-7 Individual Data Type Options

Option	Description
Data Type	Displays the unique name for the data type.
Threshold	Specifies the number of occurrences of the data type when the system generates an event. You receive an error message when you save the policy if you do not set a threshold for an enabled data type. You can specify 1 through 255. Note that the preprocessor generates one event for a detected data type per session. Note also that global threshold events are independent of individual data type events; that is, the preprocessor generates an event when the data type event threshold is reached, regardless of whether the global event threshold has been reached, and vice versa.
Destination Ports	Specifies destination ports to monitor for the data type. You can specify a single port, a comma-separated list of ports, or `any`, meaning any destination port. You receive an error message when you save the policy if you enable the rule for a data type without setting at least one port or application protocol for the data type.
Application Protocols Note that this feature requires a Control license.	Specifies up to eight application protocols to monitor for the data type. You receive an error message when you save the policy if you enable the rule for a data type without setting at least one port or application protocol for the data type. See Selecting Application Protocols to Monitor, page 28-25 for detailed instructions for selecting application protocols for data types.
Pattern	For a custom data type, the specified pattern to detect (data patterns for data types provided by Cisco are predefined). See Using Custom Data Types, page 28-27 for more information. The user interface does not display built-in patterns for predefined data types. Note that custom and predefined data patterns are system-wide.

Using Predefined Data Types

License: Protection

Each intrusion policy includes predefined data types for detecting commonly used data patterns such as credit card numbers, email addresses, U.S. phone numbers, and U.S. Social Security numbers with and without dashes. Each predefined data type is associated with a single sensitive data preprocessor rule that has a generator ID (GID) of 138. You must enable the associated sensitive data rule in the intrusion policy to enable detection, and event generation, for each data type you want to use in your policy. See Setting Rule States, page 27-19 for information on enabling rules in an intrusion policy.

To help you enable sensitive data rules, a link on the configuration page takes you to a filtered view of the Rules page that displays all predefined and custom sensitive data rules. You can also display only predefined sensitive data rules by selecting the sensitive-data rule filtering category on the Rules page. See Filtering Rules in an Intrusion Policy, page 27-9 for more information. Predefined sensitive data rules are also listed on the Rule Editor page (**Policies > Intrusion > Rule Editor**), where you can view but not edit them under the sensitive-data rule category.

The following table describes each data type and lists the corresponding preprocessor rule that you must enable to enable detection and event generation for the data type.

Table 28-8 Sensitive Data Types

Data Type	Description	Preprocessor Rule GID:SID
Credit Card Numbers	Matches Visa®, MasterCard®, Discover® and American Express® fifteen- and sixteen-digit credit card numbers, with or without their normal separating dashes or spaces; also uses the Luhn algorithm to verify credit card check digits.	138:2
Email Addresses	Matches email addresses.	138:5
U.S. Phone Numbers	Matches U.S. phone numbers adhering to the pattern `(\d{3})?\d{3}-\d{4}`.	138:6
U.S. Social Security Numbers Without Dashes	Matches 9-digit U.S. Social Security numbers that have valid 3-digit area numbers, valid 2-digit group numbers, and do not have dashes.	138:4
U.S. Social Security Numbers With Dashes	Matches 9-digit U.S. Social Security numbers that have valid 3-digit area numbers, valid 2-digit group numbers, and dashes.	138:3
Custom	Matches a user-defined data pattern in the specified traffic. See Using Custom Data Types, page 28-27 for more information.	138:>999999

To reduce false positives from 9-digit numbers other than Social Security numbers, the preprocessor uses an algorithm to validate the 3-digit area number and 2-digit group number that precede the 4-digit serial number in each Social Security number. The preprocessor validates Social Security group numbers through November 2009.

Configuring Sensitive Data Detection

License: Protection

You can modify default global settings and settings for individual data types. You must also enable the preprocessor rule for each data type you want to detect.

If you enable sensitive data preprocessor rules in your policy without enabling sensitive data detection, you are prompted to enable sensitive data detection when you save changes to your policy. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

The following table describes actions you can take on the Sensitive Data Detection page.

Table 28-9 Sensitive Data Configuration Actions

To...	You can...
modify global settings	see Table 28-6 on page 28-8 for information on the global settings you can modify.
modify data type options	click the data type name in the Targets page area. The Configuration page area updates to display the current settings for the data type. See the Individual Data Type Options table for information on the options you can modify.

Table 28-9 Sensitive Data Configuration Actions (continued)

To...	You can...
add or remove application protocols to monitor for a data type Note that this feature requires a Control license.	click inside the **Application Protocols** field, or click **Edit** next to the field. The Application Protocols pop-up window appears: • To add up to eight application protocols to monitor, select one or more application protocols from the **Available** list on the left, then click the right arrow (>) button. • To remove an application protocol, select it from the **Enabled** list on the right, then click the left arrow (<) button. Use Ctrl or Shift while clicking to select multiple application protocols. You can also click and drag to select multiple adjacent application protocols. **Note** To detect sensitive data in FTP traffic, you must add the `Ftp data` application protocol. See Special Case: Detecting Sensitive Data in FTP Traffic, page 28-26 for more information.
create a custom data type	click the + sign next to **Data Types** on the left side of the page. The Add Data Type pop-up window appears. Specify a unique data type name and the pattern you want to detect with this data type and click **OK**, or click **Cancel** to abandon your edits. See Using Custom Data Types, page 28-27 for more information.
display sensitive data preprocessor rules	click the **Configure Rules for Sensitive Data Detection** link above the Global Settings page area. A listing of all sensitive data preprocessor rules appears in a filtered display of the Rules page. Optionally, you can enable or disable any of the listed rules. Note that you must enable the sensitive data preprocessor rule for each data type that you want to use in your intrusion policy. See Setting Rule States, page 27-19 for more information. You can also configure sensitive data rules for any of the other actions available on the Rules page, such as rule suppression, rate-based attack prevention, and so on; see Tuning Intrusion Policies Using Rules, page 27-1 for more information. Click **Back** to return to the Sensitive Data Detection page.

To configure sensitive data detection:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

Step 4 You have two choices, depending on whether **Sensitive Data Detection** under **Specific Threat Detection** is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Sensitive Data Detection page appears. A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 5 You can take any of the actions described in the Sensitive Data Configuration Actions table.

Step 6 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Selecting Application Protocols to Monitor

License: Control

You can specify up to eight application protocols to monitor for each data type.

You must specify at least one application protocol or port to monitor for each data type. However, except in the case where you want to detect sensitive data in FTP traffic, Cisco recommends for the most complete coverage that you specify corresponding ports when you specify application protocols. For example, if you specify HTTP, you might also configure the well-known HTTP port 80. If a new host on your network implements HTTP, the system will monitor port 80 during the interval when it is discovering the new HTTP application protocol.

In the case where you want to detect sensitive data in FTP traffic, you must specify the `FTP data` application protocol; there is no advantage in specifying a port number. See Special Case: Detecting Sensitive Data in FTP Traffic, page 28-26 for more information.

To modify application protocols to detect sensitive data:

Admin/Intrusion Admin

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

Step 4 You have two choices, depending on whether **Sensitive Data Detection** under **Specific Threat Detection** is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Sensitive Data Detection page appears.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 5 Click the data type name under **Data Types** to select the data type you want to modify.

The Configuration area updates to display the current settings for the selected data type.

Step 6 Click inside the **Application Protocols** field, or click **Edit** next to the field.

The Application Protocols pop-up window appears.

Step 7 You have two choices:

- To add up to eight application protocols to monitor, select one or more application protocols from the **Available** list on the left, then click the right arrow (>) button.
- To remove an application protocol, select it from the **Enabled** list on the right, then click the left arrow (<) button.

Use Ctrl or Shift while clicking to select multiple application protocols. You can also click and drag to select multiple adjacent application protocols.

Note To detect sensitive data in FTP traffic, you must add the `FTP data` application protocol. See Special Case: Detecting Sensitive Data in FTP Traffic, page 28-26 for more information.

Step 8 Click **OK** to add the application protocols.

The Sensitive Data Detection page is displayed and the application protocols are updated.

Special Case: Detecting Sensitive Data in FTP Traffic

License: Control

You usually determine which traffic to monitor for sensitive data by specifying the ports to monitor or, optionally, specifying application protocols in deployments. However, specifying ports or application protocols is not sufficient for detecting sensitive data in FTP traffic. Sensitive data in FTP traffic is found in traffic for the FTP application protocol, which occurs intermittently and uses a transient port number, making it difficult to detect. To detect sensitive data in FTP traffic, you **must** include the following in your configuration:

- Specify the `FTP data` application protocol.

  Specifying the `FTP data` application protocol enables detection of sensitive data in FTP traffic. See Selecting Application Protocols to Monitor, page 28-25 for more information.

  In the special case of detecting sensitive data in FTP traffic, specifying the `FTP data` application protocol does not invoke detection; instead, it invokes the rapid processing of the FTP/Telnet processor to detect sensitive data in FTP traffic. See Decoding FTP and Telnet Traffic, page 22-18 for more information.

- Ensure that your configuration includes at least one port to monitor for sensitive data.

  Note that it is not necessary to specify an FTP port except in the unlikely case where you only want to detect sensitive data in FTP traffic. Most sensitive data configurations will include other ports such as HTTP or email ports. In the case where you do want to specify only one FTP port and no other ports to monitor, Cisco recommends that you specify the FTP command port `23`. See Configuring Sensitive Data Detection, page 28-23 or more information.

Using Custom Data Types

License: Protection

You can create and modify custom data types to detect data patterns that you specify. For example, a hospital might create a data type to protect patient numbers, or a university might create a data type to detect student numbers that have a unique numbering pattern.

Each custom date type you create also creates a single sensitive data preprocessor rule that has a generator ID (GID) of 138 and a Snort ID of 1000000 or greater, that is, a SID for a local rule. You must enable the associated sensitive data rule to enable detection, and event generation, for each custom data type you want to use in your policy. See Setting Rule States, page 27-19 for information on enabling rules in an intrusion policy.

To help you enable sensitive data rules, a link on the configuration page takes you to a filtered view of the Rules page that displays all predefined and custom sensitive data rules. You can also display custom sensitive data rules along with any local custom rules by selecting the local rule filtering category on the Rules page. See Filtering Rules in an Intrusion Policy, page 27-9 for more information. Note that custom sensitive data rules are not listed on the Rule Editor page.

Custom data types you create are added to all intrusion policies. You must enable the associated sensitive data rule in any policy that you want to use to detect and generate events for a particular custom data type.

Note that you must use the Sensitive Data Detection configuration page to create data types and their associated rules. You cannot use the rule editor to create sensitive data rules.

See the following sections for more information:

- Defining Data Patterns in Custom Data Types, page 28-27
- Configuring Custom Data Types, page 28-29
- Editing Custom Data Type Names and Detection Patterns, page 28-30

Defining Data Patterns in Custom Data Types

License: Protection

You define the data pattern for a custom data type using a simple set of regular expressions comprised of the following:

- three metacharacters
- escaped characters that allow you to use the metacharacters as literal characters
- six character classes

Metacharacters are literal characters that have special meaning within regular expressions. The following table describes the metacharacters you can use when defining a custom data pattern.

Table 28-10 Sensitive Data Pattern Metacharacters

Metacharacter	Description	Example
`?`	Matches zero or one occurrence of the preceding character or escape sequence; that is, the preceding character or escape sequence is optional.	`colou?r` matches `color` or `colour`
`{n}`	Matches the preceding character or escape sequence *n* times.	For example, `\d{2}` matches `55`, `12`, and so on; `\l{3}` matches `AbC`, `www`, and so on; `\w{3}` matches `a1B`, 25C, and so on; `x{5}` matches `xxxxx`
`\`	Allows you to use metacharacters as actual characters and is also used to specify a predefined character class. See Table 28-12 on page 28-28 for a description of the character classes you can use in sensitive data patterns.	`\?` matches a question mark, `\\` matches a backslash, `\d` matches numeric characters, and so on

You must use a backslash to escape the characters in the following table for the sensitive data preprocessor to interpret them correctly as literal characters.

Table 28-11 Escaped Sensitive Data Pattern Characters

Use this escaped character...	To represent this literal character...
\?	?
\{	{
\}	}
\\	\

The following table describes the character classes you can use when defining a custom sensitive data pattern.

Table 28-12 Sensitive Data Pattern Character Classes

Character Class	Description	Character Class Definition
\d	Matches any numeric ASCII character 0-9	0-9
\D	Matches any byte that is not a numeric ASCII character	not 0-9
\l (lowercase “ell”)	Matches any ASCII letter	a-zA-Z
\L	Matches any byte that is not an ASCII letter	not a-zA-Z

Table 28-12 Sensitive Data Pattern Character Classes (continued)

Character Class	Description	Character Class Definition
\w	Matches any ASCII alphanumeric character Note that, unlike PCRE regular expressions, this does not include an underscore (_).	a-zA-Z0-9
\W	Matches any byte that is not an ASCII alphanumeric character	not a-zA-Z0-9

The preprocessor treats characters entered directly, instead of as part of a regular expression, as literal characters. For example, the data pattern 1234 matches `1234`.

The following data pattern example, which is used in predefined sensitive data rule 138:4, uses the escaped digits character class, the multiplier and option-specifier metacharacters, and the literal dash (-) and left and right parentheses () characters to detect U.S. phone numbers:

```
(\d{3}) ?\d{3}-\d{4}
```

Exercise caution when creating custom data patterns. Consider the following alternative data pattern for detecting phone numbers which, although using valid syntax, could cause many false positives:

```
(?\d{3})? ?\d{3}-?\d{4}
```

Because the second example combines optional parentheses, optional spaces, and optional dashes, it would detect, among others, phone numbers in the following desirable patterns:

- `(555)123-4567`
- `555123-4567`
- `5551234567`

However, the second example pattern would also detect, among others, the following potentially invalid patterns, resulting in false positives:

- `(555 1234567`
- `555)123-4567`
- `555) 123-4567`

Consider finally, for illustration purposes only, an extreme example in which you create a data pattern that detects the lowercase letter `a` using a low event threshold in all destination traffic on a small company network. Such a data pattern could overwhelm your system with literally millions of events in only a few minutes.

Configuring Custom Data Types

License: Protection

You configure essentially the same data type options for custom data types that you configure for predefined data types. See Selecting Individual Data Type Options, page 28-21 for information on setting options that are common to all data types. In addition, you must also specify the name and data pattern for custom data types.

Note that creating a custom data type also creates an associated custom sensitive data preprocessing rule, which you must enable in each policy where you want to use that data type. See Setting Rule States, page 27-19 for information on enabling rules in your intrusion policy.

To create or modify a custom data type:

Admin/Intrusion Admin

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

Step 4 You have two choices, depending on whether **Sensitive Data Detection** under **Specific Threat Detection** is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Sensitive Data Detection page appears.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 5 You have the following options:

- To create a custom data type, click the + sign next to **Data Types** on the left side of the page. The Add Data Type pop-up window appears.

  Specify a unique data type name and the pattern you want to detect with this data type and click **OK**, or click **Cancel** to abandon your edits. See Editing Custom Data Type Names and Detection Patterns, page 28-30 for more information.

  The Sensitive Data Detection page appears. If you clicked **OK**, the page updates to display your changes.

- To modify any of the options that are common to predefined and custom data types, click the data type name in the **Targets** page area.

  The Configuration page area updates to display the current settings for the data type. See Configuring Sensitive Data Detection, page 28-23 for more information.

- To edit the system-wide name and data pattern for a custom data type, see Editing Custom Data Type Names and Detection Patterns, page 28-30.

- To delete a custom data type, click the delete icon ( ) next to the data type you want to remove and then click **OK**, or click **Cancel** to abandon deleting the data type.

  Note that you cannot delete a data type when the sensitive data rule for that data type is enabled in any intrusion policy. Deleting a custom data type deletes it from all intrusion policies.

Editing Custom Data Type Names and Detection Patterns

License: Protection

You can modify the system-wide name and detection pattern for custom sensitive data rules. Note that changing these settings changes them in all other policies on the system. Note also that you must reapply any applied access control policies that include intrusion policies that use custom data types that you modify.

Except for custom data type names and data patterns, all data type options are policy-specific for both custom and predefined data types. See Selecting Individual Data Type Options, page 28-21 for information on modifying options other than the name and data pattern in your custom data types.

To edit custom data type names and data patterns:

Admin/Intrusion Admin

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy.**

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

Step 4 You have two choices, depending on whether **Sensitive Data Detection** under **Specific Threat Detection** is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Sensitive Data Detection page appears.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 5 In the **Targets** page area, click the name of the custom data type you want to modify.

The page updates to show the current settings for the data type, and the **Edit Data Type Name and Pattern** link appears in the upper right of the Configuration page area.

Step 6 Click the **Edit Data Type Name and Pattern** link.

The Edit Data Type pop-up window appears.

Step 7 Modify the data type name, pattern, or both and click **OK**, or click **Cancel** to abandon your edits. See Defining Data Patterns in Custom Data Types, page 28-27 for information on specifying the data pattern.

The Sensitive Data Detection page appears. If you clicked **OK**, the page displays your changes.



Globally Limiting Intrusion Event Logging

You can use thresholds to limit the number of times the system logs and displays intrusion events. Thresholds, which you configure as part of your intrusion policy, cause the system to generate events based on how many times traffic matching a rule originates from or is targeted to a specific address or address range within a specified time period. This can prevent you from being overwhelmed with a large number of events. This feature requires a Protection license.

You can set event notification thresholds in two ways:

- You can set a global threshold across all traffic to limit how often events from a specific source or destination are logged and displayed per specified time period. For more information, see Understanding Thresholding, page 29-1 and Configuring Global Thresholds, page 29-3.
- You can set thresholds per shared object rule, standard text rule, or preprocessor rule in your intrusion policy configuration, as described in Configuring Event Thresholding, page 27-21.

Understanding Thresholding

License: Protection

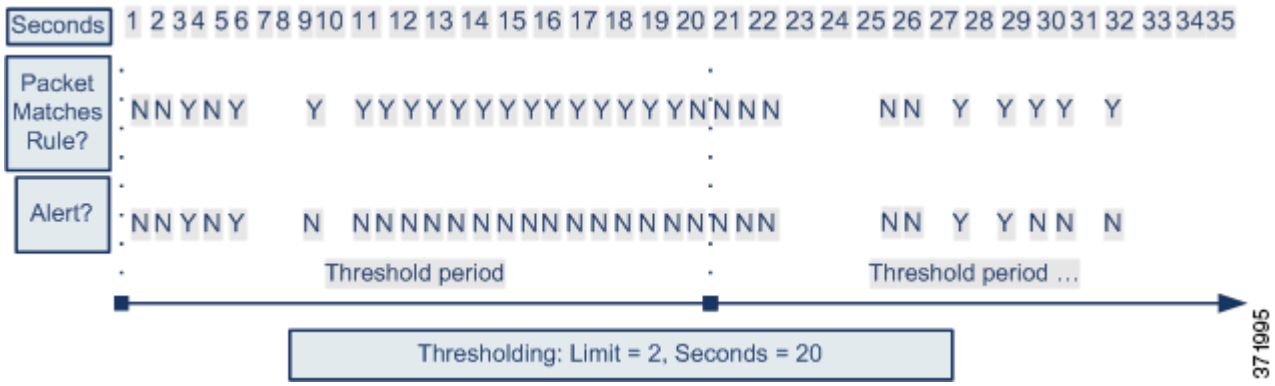
By default, every intrusion policy contains a global rule threshold. The default threshold limits event generation for each rule to one event every 60 seconds on traffic going to the same destination. This global threshold applies by default to all intrusion rules and preprocessor rules. Note that you can disable the threshold in the Advanced Settings page in an intrusion policy.

You can also override this threshold by setting individual thresholds on specific rules. For example, you might set a global limit threshold of five events every 60 seconds, but then set a specific threshold of ten events for every 60 seconds for SID 1315. All other rules generate no more than five events in each 60-second period, but the system generates up to ten events for each 60-second period for SID 1315.

For more information on setting rule-based thresholds, see Configuring Event Thresholding, page 27-21.

The following diagram shows an example where an attack is in progress for a specific rule. A global limit threshold limits event generation for each rule to two events every 20 seconds.

Note that the period starts at one second and ends at 21 seconds. After the period ends, note that the cycle starts again and the next two rule matches generate events, then the system does not generate any more events during that period.



Understanding Thresholding Options

License: Protection

Thresholding allows you to limit intrusion event generation by generating only a specific number of events in a time period, or by generating one event for a set of events. When you configure global thresholding, you must first specify the thresholding type, as described in the following table.

Table 29-1 Thresholding Options

Option	Description
Limit	Logs and displays events for the specified number of packets (specified by the count argument) that trigger the rule during the specified time period. For example, if you set the type to **Limit**, the **Count** to `10`, and the **Seconds** to `60`, and 14 packets trigger the rule, the system stops logging events for the rule after displaying the first 10 that occur within the same minute.
Threshold	Logs and displays a single event when the specified number of packets (specified by the count argument) trigger the rule during the specified time period. Note that the counter for the time restarts after you hit the threshold count of events and the system logs that event. For example, you set the type to **Threshold**, **Count** to `10`, and **Seconds** to `60`, and the rule triggers 10 times by second 33. The system generates one event, then resets the Seconds and Count counters to `0`. The rule then triggers another 10 times in the next 25 seconds. Because the counters reset to `0` at second 33, the system logs another event.
Both	Logs and displays an event once per specified time period, after the specified number (count) of packets trigger the rule. For example, if you set the type to **Both**, **Count** to `2`, and **Seconds** to `10`, the following event counts result: • If the rule is triggered once in 10 seconds, the system does not generate any events (the threshold is not met) • If the rule is triggered twice in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time) • If the rule is triggered four times in 10 seconds, the system generates one event (the threshold is met when the rule triggered the second time and following events are ignored)

Next, specify the tracking, which determines whether the event instance count is calculated per source or destination IP address. Finally, specify the number of instances and time period that define the threshold.

Table 29-2 Thresholding Instance/Time Options

Option	Description
Count	The number of event instances per specified time period per tracking IP address or address range required to meet the threshold.
Seconds	The number of seconds that elapse before the count resets. If you set the threshold type to **Limit**, the tracking to **Source**, **Count** to 10, and **Seconds** to 10, the system logs and displays the first 10 events that occur in 10 seconds from a given source port. If only seven events occur in the first 10 seconds, the system logs and displays those, if 40 events occur in the first 10 seconds, the system logs and displays 10, then begins counting again when the 10-second time period elapses.

Configuring Global Thresholds

License: Protection

You can set a global threshold to manage the number of events generated by each rule over a period of time. When you set a global threshold, that threshold applies for each rule that does not have an overriding specific threshold. For more information on configuring thresholds, see Understanding Thresholding, page 29-1.

A global threshold is configured on your system by default. The default values are as follows:

- **Type** — Limit
- **Track By** — Destination
- **Count** — 1
- **Seconds** — 60

To configure global thresholding:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy**.

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

Step 4 You have two choices, depending on whether **Global Rule Thresholding** under **Intrusion Rule Thresholds** is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Global Rule Thresholding page appears. A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 5 From the **Type** radio buttons, select the type of threshold that will apply over the time specified by the seconds argument. See the Thresholding Options table for more information:

- Select **Limit** to log and display an event for each packet that triggers the rule until the limit specified by the count argument is exceeded.
- Select **Threshold** to log and display a single event for each packet that triggers the rule and represents either the instance that matches the threshold set by the count argument or is a multiple of the threshold.
- Select **Both** to log and display a single event after the number of packets specified by the count argument trigger the rule.

Step 6 Select the tracking method from the **Track By** radio buttons:

- Select **Source** to identify rule matches in traffic coming from a particular source IP address or addresses.
- Select **Destination** to identify rule matches in traffic going to a particular destination IP address.

Step 7 In the **Count** field:

- For a **Limit** threshold, specify the number of event instances per specified time period per tracking IP address required to meet the threshold.
- For a **Threshold** threshold, specify the number of rule matches you want to use as your threshold.

Step 8 In the **Seconds** field:

- For a **Limit** threshold, specify the number of seconds that make up the time period when attacks are tracked.
- For a **Threshold** threshold, specify the number of seconds that elapse before the count resets. Note that the count resets if the number of rule matches indicated by the **Count** field occur before the number of seconds indicated elapse.

Step 9 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Disabling the Global Threshold

License: Protection

By default, a global limit threshold limits the number of events on traffic going to a destination to one event per 60 seconds. You can disable global thresholding in the highest policy layer if you want to threshold events for specific rules and not apply thresholding to every rule by default.

To disable global thresholding:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy**.

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Settings** in the navigation panel on the left.

The Settings page appears.

Step 4 Under **Intrusion Rule Thresholds**, disable **Global Rule Thresholding**.

Step 5 Save your policy, continue editing, discard your changes, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.



Understanding and Writing Intrusion Rules

An *intrusion rule* is a specified set of keywords and arguments that detects attempts to exploit vulnerabilities on your network by analyzing network traffic to check if it matches the criteria in the rule. The system compares packets against the conditions specified in each rule and, if the packet data matches all the conditions specified in a rule, the rule triggers. If a rule is an *alert rule*, it generates an intrusion event. If it is a *pass rule*, it ignores the traffic. You can view and evaluate intrusion events from the ASA FirePOWER module interface.

Caution Make sure you use a controlled network environment to test any intrusion rules that you write before you use the rules in a production environment. Poorly written intrusion rules may seriously affect the performance of the system.

Note the following:

- For a *drop* rule in an inline deployment, the system drops the packet and generates an event. For more information on drop rules, see Setting Rule States, page 27-19.
- Cisco provides two types of intrusion rules: shared object rules and standard text rules. The Cisco Vulnerability Research Team (VRT) can use shared object rules to detect attacks against vulnerabilities in ways that traditional standard text rules cannot. You cannot create shared object rules. When you write your own intrusion rule, you create a standard text rule.

You can write custom standard text rules to tune the types of events you are likely to see. Note that while this documentation sometimes discusses rules targeted to detect specific exploits, the most successful rules target traffic that may attempt to exploit known vulnerabilities rather than specific known exploits. By writing rules and specifying the rule's event message, you can more easily identify traffic that indicates attacks and policy evasions. For more information about evaluating events, see Viewing Events, page 37-1.

When you enable a custom standard text rule in a custom intrusion policy, keep in mind that some rule keywords and arguments require that traffic first be decoded or preprocessed in a certain way. This chapter explains the options you must configure in your network analysis policy, which governs preprocessing. Note that if you disable a required preprocessor, the system automatically uses it with its current settings, although the preprocessor remains disabled in the network analysis policy user interface.

Note Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task. For more information, see Limitations of Custom Policies, page 18-11.

See the following sections for more information:

- Understanding Rule Anatomy, page 30-2 describes the components, including the rule header and rule options, that make up a valid standard text rule.
- Understanding Rule Headers, page 30-3 provides a detailed description of the parts of a rule header.
- Understanding Keywords and Arguments in Rules, page 30-9 explains the usage and syntax of the intrusion rule keywords available in the ASA FirePOWER module.
- Constructing a Rule, page 30-101 explains how to build a new rule using the rule editor.
- Filtering Rules on the Rule Editor Page, page 30-105 explains how to display a subset of rules to help you find specific rules.

Understanding Rule Anatomy

License: Protection

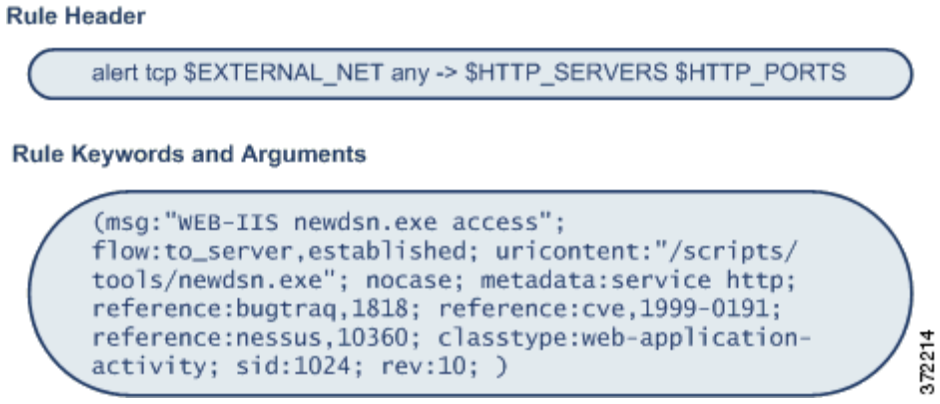
All standard text rules contain two logical sections: the rule header and the rule options. The rule header contains:

- the rule's action or type
- the protocol
- the source and destination IP addresses and netmasks
- direction indicators showing the flow of traffic from source to destination
- the source and destination ports

The rule options section contains:

- event messages
- keywords and their parameters and arguments
- patterns that a packet’s payload must match to trigger the rule
- specifications of which parts of the packet the rules engine should inspect

The following diagram illustrates the parts of a rule:



Note that the options section of a rule is the section enclosed in parentheses. The rule editor provides an easy-to-use interface to help you build standard text rules.

Understanding Rule Headers

License: Protection

Every standard text rule and shared object rule has a rule header containing parameters and arguments. The following illustrates parts of a rule header:



The following table describes each part of the rule header shown above.

Table 30-1 Rule Header Values

Rule Header Component	Example Value	This Value...
Action	`alert`	Generates an intrusion event when triggered.
Protocol	`tcp`	Tests TCP traffic only.
Source IP Address	`$EXTERNAL_NET`	Tests traffic coming from any host that is not on your internal network.
Source Ports	`any`	Tests traffic coming from any port on the originating host.
Operator	`->`	Tests external traffic (destined for the web servers on your network).
Destination IP Address	`$HTTP_SERVERS`	Tests traffic to be delivered to any host specified as a web server on your internal network.
Destination Ports	`$HTTP_PORTS`	Tests traffic delivered to an HTTP port on your internal network.

Note The previous example uses default variables, as do most intrusion rules. See Working with Variable Sets, page 2-13 for more information about variables, what they mean, and how to configure them.

See the following sections for more information about rule header parameters:

- Specifying Rule Actions, page 30-4 describes rule types and explains how to specify the action that occurs when the rule triggers.
- Specifying Protocols, page 30-4 explains how to define the traffic protocol for traffic that the rule should test.
- Specifying IP Addresses In Intrusion Rules, page 30-5 explains how to define the individual IP addresses and IP address blocks in the rule header.
- Defining Ports in Intrusion Rules, page 30-8 explains how to define the individual ports and port ranges in the rule header.

- Specifying Direction, page 30-9 describes the available operators and explains how to specify the direction traffic must be traveling to be tested by the rule.

Specifying Rule Actions

License: Protection

Each rule header includes a parameter that specifies the action the system takes when a packet triggers a rule. Rules with the action set to *alert* generate an intrusion event against the packet that triggered the rule and log the details of that packet. Rules with the action set to *pass* do not generate an event against, or log the details of, the packet that triggered the rule.

Note In an inline deployment, rules with the rule state set to *Drop and Generate Events* generate an intrusion event against the packet that triggered the rule. Also, if you apply a drop rule in a passive deployment, the rule acts as an alert rule. For more information on drop rules, see Setting Rule States, page 27-19.

By default, pass rules override alert rules. You can create pass rules to prevent packets that meet criteria defined in the pass rule from triggering the alert rule in specific situations, rather than disabling the alert rule. For example, you might want a rule that looks for attempts to log into an FTP server as the user "anonymous" to remain active. However, if your network has one or more legitimate anonymous FTP servers, you could write and activate a pass rule that specifies that, for those specific servers, anonymous users do not trigger the original rule.

Within the rule editor, you select the rule type from the **Action** list. For more information about the procedures you use to build a rule header using the rule editor, see Constructing a Rule, page 30-101.

Specifying Protocols

License: Protection

In each rule header, you must specify the protocol of the traffic the rule inspects. You can specify the following network protocols for analysis:

- ICMP (Internet Control Message Protocol)
- IP (Internet Protocol)

Note The system ignores port definitions in an intrusion rule header when the protocol is set to `ip`. For more information, see Defining Ports in Intrusion Rules, page 30-8.

- TCP (Transmission Control Protocol)
- UDP (User Datagram Protocol)

Use **IP** as the protocol type to examine all protocols assigned by IANA, including TCP, UDP, ICMP, IGMP, and many more. See http://www.iana.org/assignments/protocol-numbers for a full list of IANA-assigned protocols.

Note You cannot currently write rules that match patterns in the next header (for example, the TCP header) in an IP payload. Instead, content matches begin with the last decoded protocol. As a workaround, you can match patterns in TCP headers by using rule options.

Within the rule editor, you select the protocol type from the **Protocol** list. See Constructing a Rule, page 30-101 for more information about the procedures you use to build a rule header using the rule editor.

Specifying IP Addresses In Intrusion Rules

License: Protection

Restricting packet inspection to the packets originating from specific IP addresses or destined to a specific IP address reduces the amount of packet inspection the system must perform. This also reduces false positives by making the rule more specific and removing the possibility of the rule triggering against packets whose source and destination IP addresses do not indicate suspicious behavior.

Tip The system recognizes only IP addresses and does not accept host names for source or destination IP addresses.

Within the rule editor, you specify source and destination IP addresses in the **Source IPs** and **Destination IPs** fields. See Constructing a Rule, page 30-101 for more information about the procedures you use to build a rule header using the rule editor.

When writing standard text rules, you can specify IPv4 and IPv6 addresses in a variety of ways, depending on your needs. You can specify a single IP address, `any`, IP address lists, CIDR notation, prefix lengths, a network variable, or a network object or network object group. Additionally, you can indicate that you want to exclude a specific IP address or set of IP addresses. When specifying IPv6 addresses, you can use any addressing convention defined in RFC 4291.

The following table summarizes the various ways you can specify source and destination IP addresses.

Table 30-2 Source/Destination IP Address Syntax

To Specify...	Use...	Example
any IP address	`any`	`any`
a specific IP address	the IP address Note that you would not mix IPv4 and IPv6 source and destination addresses in the same rule.	`192.168.1.1` `2001:db8::abcd`
a list of IP addresses	brackets (`[]`) to enclose the IP addresses and commas to separate them	`[192.168.1.1,192.168.1.15]` `[2001:db8::b3ff, 2001:db8::0202]`
a block of IP addresses	IPv4 CIDR block or IPv6 address prefix notation	`192.168.1.0/24` `2001:db8::/32`
anything except a specific IP address or set of addresses	the `!` character before the IP address or addresses you want to negate	`!192.168.1.15` `!2001:db8::0202:b3ff:fe1e`
anything in a block of IP addresses except one or more specific IP addresses	a block of addresses followed by a list of negated addresses or blocks	`[10.0.0/8, !10.2.3.4, !10.1.0.0/16]` `[2001:db8::/32, !2001:db8::8329, !2001:db8::0202]`

Table 30-2 Source/Destination IP Address Syntax (continued)

To Specify...	Use...	Example
IP addresses defined by a network variable	the variable name, in uppercase letters, preceded by `$` Note that preprocessor rules can trigger events regardless of the hosts defined by network variables used in intrusion rules. See Working with Variable Sets, page 2-13 for more information.	`$HOME_NET`
all IP addresses except addresses defined by an IP address variable	the variable name, in uppercase letters, preceded by `!$` See Excluding IP Addresses in Intrusion Rules, page 30-7 for more information.	`!$HOME_NET`
IP addresses defined by a network object or network object group	the object or group name using the format `!{object_name}`. See Working with Network Objects, page 2-3 for more information.	`${192.168sub16}`
all IP addresses except addresses defined by a network object or network object group	the object or group name, in curly braces (`{}`), preceded by `!$`. See Working with Network Objects, page 2-3 for more information.	`!${192.168sub16}`

See the following sections for more in-depth information about the syntax you can use to specify source and destination IP addresses, and for information about using variables to specify IP addresses:

- IP Address Conventions, page 1-4.
- Working with Variable Sets, page 2-13
- Specifying Any IP Address, page 30-6
- Specifying Multiple IP Addresses, page 30-6
- Specifying Network Objects, page 30-7
- Excluding IP Addresses in Intrusion Rules, page 30-7

Specifying Any IP Address

License: Protection

You can specify the word `any` as a rule source or destination IP address to indicate any IPv4 or IPv6 address.

For example, the following rule uses the argument **`any`** in the **Source IPs** and **Destination IPs** fields and evaluates packets with any IPv4 or IPv6 source or destination address:

```
alert tcp any any -> any any
```

You can also specify `::` to indicate any IPv6 address.

Specifying Multiple IP Addresses

License: Protection

You can list individual IP addresses by separating the IP addresses with commas and, optionally, by surrounding non-negated lists with brackets, as shown in the following example:

```
[192.168.1.100,192.168.1.103,192.168.1.105]
```

You can list IPv4 and IPv6 addresses alone or in any combination, as shown in the following example:

```
[192.168.1.100,2001:db8::1234,192.168.1.105]
```

Note that surrounding an IP address list with brackets, which was required in earlier software releases, is not required. Note also that, optionally, you can enter lists with a space before or after each comma.

Note You must surround negated lists with brackets. See Excluding IP Addresses in Intrusion Rules, page 30-7 for more information.

You can also use IPv4 Classless Inter-Domain Routing (CIDR) notation or IPv6 prefix lengths to specify address blocks. For example:

- 192.168.1.0/24 specifies the IPv4 addresses in the 192.168.1.0 network with a subnet mask of 255.255.255.0, that is, 192.168.1.0 through 192.168.1.255. For more information, see IP Address Conventions, page 1-4.
- 2001:db8::/32 specifies the IPv6 addresses in the 2001:db8:: network with a prefix length of 32 bits, that is, 2001:db8:: through 2001:db8:ffff:ffff:ffff:ffff:ffff:ffff.

Tip If you need to specify a block of IP addresses but cannot express it using CIDR or prefix length notation alone, you can use CIDR blocks and prefix lengths in an IP address list.

Specifying Network Objects

License: Protection

You can specify a network object or network object group using the syntax:

```
${object_name | group_name}
```

where:

- *object_name* is the name of a network object
- *group_name* is the name of a network object group

See Working with Network Objects, page 2-3 for information on creating network objects and network object groups.

Consider the case where you have created a network object named `192.168sub16` and a network object group named `all_subnets`. You could specify the following to identify IP addresses using the network object:

```
${192.168sub16}
```

and you could specify the following to use the network object group:

```
${all_subnets}
```

You can also use negation with network objects and network object groups. For example:

```
!${192.168sub16}
```

See Excluding IP Addresses in Intrusion Rules, page 30-7 for more information.

Excluding IP Addresses in Intrusion Rules

License: Protection

You can use an exclamation point (`!`) to negate a specified IP address. That is, you can match any IP address with the exception of the specified IP address or addresses. For example, `!192.168.1.1` specifies any IP address other than 192.168.1.1, and `!2001:db8:ca2e::fa4c` specifies any IP address other than 2001:db8:ca2e::fa4c.

To negate a list of IP addresses, place `!` before a bracketed list of IP addresses. For example, `![192.168.1.1,192.168.1.5]` would define any IP address other than 192.168.1.1 or 192.168.1.5.

Note You must use brackets to negate a list of IP addresses.

Be careful when using the negation character with IP address lists. For example, if you use `[!192.168.1.1,!192.168.1.5]` to match any address that is not 192.168.1.1 or 192.168.1.5, the system interprets this syntax as "anything that is not 192.168.1.1, **or** anything that is not 192.168.1.5."

Because 192.168.1.5 is not 192.168.1.1, and 192.168.1.1 is not 192.168.1.5, both IP addresses match the IP address value of `[!192.168.1.1,!192.168.1.5]`, and it is essentially the same as using "`any`."

Instead, use `![192.168.1.1,192.168.1.5]`. The system interprets this as "**not** 192.168.1.1 **and not** 192.168.1.5," which matches any IP address other than those listed between brackets.

Note that you cannot logically use negation with `any` which, if negated, would indicate no address.

Defining Ports in Intrusion Rules

License: Protection

Within the rule editor, you specify source and destination ports in the **Source Port** and **Destination Port** fields. See Constructing a Rule, page 30-101 for more information about the procedures you use to build a rule header using the rule editor.

The ASA FirePOWER module uses a specific type of syntax to define the port numbers used in rule headers.

Note The system ignores port definitions in an intrusion rule header when the protocol is set to `ip`. For more information, see Specifying Protocols, page 30-4.

You can list ports by separating the ports with commas, as shown in the following example:

```
80, 8080, 8138, 8600-9000, !8650-8675
```

Optionally, the following example shows how you can surround a port list with brackets, which was required in previous software versions but is no longer required:

```
[80, 8080, 8138, 8600-9000, !8650-8675]
```

Note that you **must** surround negated port lists in brackets, as shown in the following example:

```
![20, 22, 23]
```

Note also that a list of source or destination ports in an intrusion rule can include a maximum of 64 characters.

The following table summarizes the syntax you can use:

Table 30-3 Source/Destination Port Syntax

To Specify...	Use	Example
any port	`any`	`any`
a specific port	the port number	`80`
a range of ports	a dash between the first and last port number in the range	`80-443`
all ports less than or equal to a specific port	a dash before the port number	`-21`
all ports greater than or equal to a specific port	a dash after the port number	`80-`
all ports except a specific port or range of ports	the `!` character before the port, port list, or range of ports you want to negate Note that you can logically use negation with all port designations except *any*, which if negated would indicate *no port*.	`!20`
all ports defined by a port variable	the variable name, in uppercase letter, preceded by `$` See Working with Port Variables, page 2-24 for more information.	`$HTTP_PORTS`
all ports except ports defined by a port variable	the variable name, in uppercase letter, preceded by `!$`	`!$HTTP_PORTS`

Specifying Direction

License: Protection

Within the rule header, you can specify the direction that the packet must travel for the rule to inspect it. The following table describes these options.

Table 30-4 Directional Options in Rule Headers

Use...	To Test...
Directional	only traffic from the specified source IP address to the specified destination IP address
Bidirectional	all traffic traveling between the specified source and destination IP addresses

See Constructing a Rule, page 30-101 for more information about the procedures you use to build a rule header using the rule editor.

Understanding Keywords and Arguments in Rules

License: Protection

Using the rules language, you can specify the behavior of a rule by combining keywords. Keywords and their associated values (called *arguments*) dictate how the system evaluates packets and packet-related values that the rules engine tests. The ASA FirePOWER module currently supports keywords that allow you to perform inspection functions, such as content matching, protocol-specific pattern matching, and state-specific matching. You can define up to 100 arguments per keyword, and combine any number of compatible keywords to create highly specific rules. This helps decrease the chance of false positives and false negatives and focus the intrusion information you receive.

Note that you can also use adaptive profiles to dynamically adapt active rule processing for specific packets based on rule metadata and host information. For more information, see Tuning Preprocessing in Passive Deployments, page 25-1.

See the following sections for more information:

- Defining Intrusion Event Details, page 30-11 describes the syntax and use of keywords that allow you to define the event's message, priority information, and references to external information about the exploit the rule detects.
- Searching for Content Matches, page 30-15 describes how to use the `content` or `protected_content` keywords to test the content of the packet payload.
- Constraining Content Matches, page 30-17 describes how to use modifying keywords for the `content` or `protected_content` keywords.
- Replacing Content in Inline Deployments, page 30-29 describes how to use the `replace` keyword in inline deployments to replace specified content of equal length.
- Using Byte_Jump and Byte_Test, page 30-30 describes how to use the `byte_jump` and `byte_test` keywords to calculate where in a packet the rules engine should begin testing for a content match, and which bytes it should evaluate.
- Searching for Content Using PCRE, page 30-35 describes how to use the `pcre` keyword to use Perl-compatible regular expressions in rules.
- Adding Metadata to a Rule, page 30-43 describes how to use the `metadata` keyword to add information to a rule.
- Inspecting IP Header Values, page 30-45 describes the syntax and use of keywords that test values in the packet's IP header.
- Inspecting ICMP Header Values, page 30-48 describes the syntax and use of keywords that test values in the packet's ICMP header.
- Inspecting TCP Header Values and Stream Size, page 30-50 describes the syntax and use of keywords that test values in the packet's TCP header.
- Enabling and Disabling TCP Stream Reassembly, page 30-54 describes how to enable and disable stream reassembly for a single connection when inspected traffic on the connection matches the conditions of the rule.
- Extracting SSL Information from a Session, page 30-54 describes the use and syntax of keywords that extract version and state information from encrypted traffic.
- Reading Packet Data into Keyword Arguments, page 30-82 describes how to read a value from a packet into a variable that you can use later in the same rule to specify the value for arguments in certain other keywords.
- Inspecting Application Layer Protocol Values, page 30-56 describes the use and syntax of keywords that test application layer protocol properties.
- Inspecting Packet Characteristics, page 30-79 describes the use and syntax of the `dsize`, `sameIP`, `isdataat`, `fragoffset`, and `cvs` keywords.
- Initiating Active Responses with Rule Keywords, page 30-84 explains how to use the `resp` keyword to actively close TCP connections or UDP sessions, the `react` keyword to send an HTML page and then actively close TCP connections, and the `config response` command to specify the active response interface and the number of TCP resets to attempt in a passive deployment.
- Filtering Events, page 30-87 describes how to prevent a rule from triggering an event unless a specified number packets meet the rule's detection criteria within a specified time.

- Evaluating Post-Attack Traffic, page 30-88 describes how to log additional traffic for the host or session.
- Detecting Attacks That Span Multiple Packets, page 30-89 describes how to assign state names to packets from attacks that span multiple packets in a single session, then analyze and alert on packets according to their state.
- Generating Events on the HTTP Encoding Type and Location, page 30-94 describes how to generate events on the type of encoding in an HTTP request or response URI, header, or cookie, including set-cookies, before normalization.
- Detecting File Types and Versions, page 30-96 describes how to point to a specific file type or file version using the `file_type` or `file_group` keyword.
- Pointing to a Specific Payload Type, page 30-98 describes how to point to the beginning of the HTTP response entity body, SMTP payload, or encoded email attachment.
- Pointing to the Beginning of the Packet Payload, page 30-99 describes how to point to the beginning of the packet payload.
- Decoding and Inspecting Base64 Data, page 30-99 describes how you can use the `base64_decode` and `base64_data` keywords to decode and inspect Base64 data, especially in HTTP requests.

Defining Intrusion Event Details

License: Protection

As you construct a standard text rule, you can include contextual information that describes the vulnerability that the rule detects attempts to exploit. You can also include external references to vulnerability databases and define the priority that the event holds in your organization. When analysts see the event, they then have information about the priority, exploit, and known mitigation readily available.

See the following sections for more information about event-related keywords:

- Defining the Event Message, page 30-11
- Defining the Event Priority, page 30-12
- Defining the Intrusion Event Classification, page 30-12
- Defining the Event Reference, page 30-14

Defining the Event Message

License: Protection

You can specify meaningful text that appears as a message when the rule triggers. The message gives immediate insight into the nature of the vulnerability that the rule detects attempts to exploit. You can use any printable standard ASCII characters except curly braces (`{}`). The system strips quotes that completely surround the message.

Tip You must specify a rule message. Also, the message cannot consist of white space only, one or more quotation marks only, one or more apostrophes only, or any combination of just white space, quotation marks, or apostrophes.

To define the event message in the rule editor, enter the event message in the **Message** field. See Constructing a Rule, page 30-101 for more information about using the rule editor to build rules.

Defining the Event Priority

License: Protection

By default, the priority of a rule derives from the event classification for the rule. However, you can override the classification priority for a rule by adding the `priority` keyword to the rule.

To specify a priority using the rule editor, select **priority** from the **Detection Options** list, and select **high**, **medium**, or **low** from the drop-down list. For example, to assign a **high** priority for a rule that detects web application attacks, add the `priority` keyword to the rule and select **high** as the priority. See Constructing a Rule, page 30-101 for more information about using the rule editor to build rules.

Defining the Intrusion Event Classification

License: Protection

For each rule, you can specify an attack classification that appears in the packet display of the event. The following table lists the name and number for each classification.

Table 30-5 Rule Classifications

Number	Classification Name	Description
1	not-suspicious	Not Suspicious Traffic
2	unknown	Unknown Traffic
3	bad-unknown	Potentially Bad Traffic
4	attempted-recon	Attempted Information Leak
5	successful-recon-limited	Information Leak
6	successful-recon-largescale	Large Scale Information Leak
7	attempted-dos	Attempted Denial of Service
8	successful-dos	Denial of Service
9	attempted-user	Attempted User Privilege Gain
10	unsuccessful-user	Unsuccessful User Privilege Gain
11	successful-user	Successful User Privilege Gain
12	attempted-admin	Attempted Administrator Privilege Gain
13	successful-admin	Successful Administrator Privilege Gain
14	rpc-portmap-decode	Decode of an RPC Query
15	shellcode-detect	Executable Code was Detected
16	string-detect	A Suspicious String was Detected
17	suspicious-filename-detect	A Suspicious Filename was Detected
18	suspicious-login	An Attempted Login Using a Suspicious Username was Detected
19	system-call-detect	A System Call was Detected
20	tcp-connection	A TCP Connection was Detected
21	trojan-activity	A Network Trojan was Detected
22	unusual-client-port-connection	A Client was Using an Unusual Port

Table 30-5 Rule Classifications (continued)

Number	Classification Name	Description
23	network-scan	Detection of a Network Scan
24	denial-of-service	Detection of a Denial of Service Attack
25	non-standard-protocol	Detection of a Non-Standard Protocol or Event
26	protocol-command-decode	Generic Protocol Command Decode
27	web-application-activity	Access to a Potentially Vulnerable Web Application
28	web-application-attack	Web Application Attack
29	misc-activity	Misc Activity
30	misc-attack	Misc Attack
31	icmp-event	Generic ICMP Event
32	inappropriate-content	Inappropriate Content was Detected
33	policy-violation	Potential Corporate Privacy Violation
34	default-login-attempt	Attempt to Login By a Default Username and Password
35	sdf	Sensitive Data
36	malware-cnc	Known malware command and control traffic
37	client-side-exploit	Known client side exploit attempt
38	file-format	Known malicious file or file based exploit

To specify a classification in the rule editor, select a classification from the **Classification** list. See Writing New Rules, page 30-101 for more information on the rule editor.

Adding Custom Classifications

License: Protection

If you want more customized content for the packet display description of the events generated by a rule you define, create a custom classification.

To add classifications to the Classification list:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy > Rule Editor**.

The Rule Editor page appears.

Step 2 Click **Create Rule**.

The Create Rule page appears.

Step 3 Under the **Classification** drop-down list, click **Edit Classifications**.

A pop-up window appears.

Step 4 Type the name of the classification in the **Classification Name** field.

You can use up to 255 alphanumeric characters, but the page is difficult to read if you use more than 40 characters. The following characters are not supported: `<>()\'"&$;` and the space character.

Step 5 Type a description of the classification in the **Classification Description** field.

You can use up to 255 alphanumeric characters and spaces. The following characters are not supported: `<>()\'"&$;`

Step 6 Select a priority from the **Priority** list.

You can select **high**, **medium**, or **low**.

Step 7 Click **Add**.

The new classification is added to the list and becomes available for use in the rule editor.

Step 8 Click **Done**.

Defining the Event Reference

License: Protection

You can use the `reference` keyword to add references to external web sites and additional information about the event. Adding a reference provides analysts with an immediately available resource to help them identify why the packet triggered a rule. The following table lists some of the external systems that can provide data on known exploits and attacks.

Table 30-6 External Attack Identification Systems

System ID	Description	Example ID
`bugtraq`	Bugtraq page	`8550`
`cve`	Common Vulnerabilities and Exposure page	`CAN-2003-0702`
`mcafee`	McAfee page	`98574`
`url`	Website reference	`www.example.com?exploit=14`
`msb`	Microsoft security bulletin	`MS11-082`
`nessus`	Nessus page	`10039`
`secure-url`	Secure Website Reference (https://...)	`intranet/exploits/exploit=14` Note that you can use `secure-url` with any secure website.

To specify a reference using the rule editor, select **reference** from the **Detection Options** list, and enter a value in the corresponding field as follows:

`id_system``,`*`id`*

where *`id_system`* is the system being used as a prefix, and *`id`* is the Bugtraq ID, CVE number, Arachnids ID, or URL (without `http://`).

For example, to specify the authentication bypass vulnerability on Microsoft Commerce Server 2002 servers documented in Bugtraq ID 17134, enter the following in the **reference** field:

```
bugtraq,17134
```

Note the following when adding references to a rule:

- Do not use a space after the comma.
- Do not use uppercase letters in the system ID.

See Constructing a Rule, page 30-101 for more information about using the rule editor to build rules.

Searching for Content Matches

License: Protection

Use the `content` keyword or the `protected_content` keyword to specify content that you want to detect in a packet. See the following sections for more information:

- Using the content Keyword, page 30-15
- Using the protected_content Keyword, page 30-15
- Configuring Content Matching, page 30-16

Using the content Keyword

When you use the `content` keyword, the rules engine searches the packet payload or stream for that string. For example, if you enter `/bin/sh` as the value for one of the `content` keywords, the rules engine searches the packet payload for the string `/bin/sh`.

Match content using either an ASCII string, hexadecimal content (binary byte code), or a combination of both. Surround hexadecimal content with pipe characters (|) in the keyword value. For example, you can mix hexadecimal content and ASCII content using something that looks like `|90C8 C0FF FFFF|/bin/sh`.

You can specify multiple content matches in a single rule. To do this, use additional instances of the `content` keyword. For each content match, you can indicate that content matches must be found in the packet payload or stream for the rule to trigger.

Using the protected_content Keyword

The `protected_content` keyword allows you to encode your search content string before configuring the rule argument. The original rule author uses a hash function (SHA-512, SHA-256, or MD5) to encode the string before configuring the keyword.

When you use the `protected_content` keyword instead of the `content` keyword, there is no change to how the rules engine searches the packet payload or stream for that string and most of the keyword options function as expected. The following table summarizes the exceptions, where the `protected_content` keyword options differ from the `content` keyword options.

Table 30-7 protected_content Option Exceptions

Option	Description
Hash Type	New option for the `protected_content` rule keyword. For more information, see Hash Type, page 30-18.
Case Insensitive	Not supported
Within	Not supported
Depth	Not supported
Length	New option for the `protected_content` rule keyword. For more information, see Length, page 30-21.
Use Fast Pattern Matcher	Not supported
Fast Pattern Matcher Only	Not supported
Fast Pattern Matcher Offset and Length	Not supported

Cisco recommends that you include at least one `content` keyword in rules that include a `protected_content` keyword to ensure that the rules engine uses the fast pattern matcher, which increases processing speed and improves performance. Position the content `keyword` before the `protected_content` keyword in the rule. Note that the rules engine uses the fast pattern matcher when a rule includes at least one `content` keyword, regardless of whether you enable the `content` keyword Use Fast Pattern Matcher argument.

Configuring Content Matching

You should almost always follow a `content` or `protected_content` keyword by modifiers that indicate where the content should be searched for, whether the search is case sensitive, and other options. See Constraining Content Matches for more information about modifiers to the `content` and `protected_content` keywords.

Note that all content matches must be true for the rule to trigger an event, that is, each content match has an AND relationship with the others.

Note also that, in an inline deployment, you can set up rules that match malicious content and then replace it with your own text string of equal length. See Replacing Content in Inline Deployments, page 30-29 for more information.

To enter content to be matched:

Step 1 In the `content` field, type the content you want to find (for example, `|90C8 C0FF FFFF|/bin/sh`).

If you want to search for any content that is **not** the specified content, select the **Not** check box.

Caution You may invalidate your intrusion policy if you create a rule that includes only one `content` keyword and that keyword has the **Not** option selected. For more information, see Not, page 30-19.

Step 2 Optionally, add additional keywords that modify the `content` keyword or add constraints for the keyword.

For more information on other keywords, see Understanding Keywords and Arguments in Rules, page 30-9. For more information on constraining the `content` keyword, see Constraining Content Matches, page 30-17.

Step 3 Continue with creating or editing the rule.

See Writing New Rules, page 30-101 or Modifying Existing Rules, page 30-103 for more information.

To enter protected content to be matched:

Step 1 Using a SHA-512, SHA-256, or MD5 hash generator, encode the content you want to find (for example, run the string `Sample1` through a SHA-512 hash generator).

The generator outputs a hash for your string.

Step 2 In the `protected_content` field, type the hash you generated in step 1 (for example, `B20AABAF59605118593404BD42FE69BD8D6506EE7F1A71CE6BB470B1DF848C814BC5DBEC2081999F15691A71FAECA5FBA4A3F8B8AB56B7F04585DA6D73E5DD15`).

If you want to search for any content that is **not** the specified content, select the **Not** check box.

Caution You may invalidate your intrusion policy if you create a rule that includes only one `protected_content` keyword and that keyword has the **Not** option selected. For more information, see Not, page 30-19.

Step 3 From the **Hash Type** drop-down list, select the hash function you used in step 1 (for example, **SHA-512**). Note that the number of bits in the hash entered in step 2 **must** match the hash type or the system does not save the rule. For more information, see Hash Type, page 30-18.

Tip If you select the Cisco-set **Default**, the system assumes SHA-512 as the hash function.

Step 4 Type a value in the required **Length** field. The value **must** correspond with the length of the original, unhashed string you want to find (for example, the string `Sample1` from step 2 has the length `7`).

For more information, see Length, page 30-21.

Step 5 Type a value in either the **Offset** or **Distance** field. You cannot mix the **Offset** and **Distance** options within a single keyword configuration.

For more information, see Using Search Location Options in the protected_content Keyword, page 30-22.

Step 6 Optionally, add additional constraining options that modify the `protected_content` keyword.

For more information, see Constraining Content Matches, page 30-17.

Step 7 Optionally, add additional keywords that modify the `protected_content` keyword.

For more information, see Understanding Keywords and Arguments in Rules, page 30-9.

Step 8 Continue with creating or editing the rule.

See Writing New Rules, page 30-101 or Modifying Existing Rules, page 30-103 for more information.

Constraining Content Matches

License: Protection

You can constrain the location and case-sensitivity of content searches with parameters that modify the `content` or `protected_content` keyword. Configure options that modify the `content` or `protected_content` keyword to specify the content for which you want to search.

For more information, see the following sections:

- Case Insensitive, page 30-18
- Hash Type, page 30-18
- Raw Data, page 30-19
- Not, page 30-19
- Search Location Options, page 30-20
- HTTP Content Options, page 30-23
- Use Fast Pattern Matcher, page 30-26

Case Insensitive

License: Protection

Note This option is **not** supported when configuring the `protected_content` keyword. For more information, see Using the protected_content Keyword, page 30-15.

You can instruct the rules engine to ignore case when searching for content matches in ASCII strings. To make your search case-insensitive, check **Case Insensitive** when specifying a content search.

To specify Case Insensitive when doing a content search:

Step 1 Select **Case Insensitive** for the `content` keyword you are adding.

Step 2 Continue with creating or editing the rule.

See Constraining Content Matches, Searching for Content Matches, page 30-15, Writing New Rules, page 30-101 or Modifying Existing Rules, page 30-103 for more information.

Hash Type

License: Protection

Note This option is **only** configurable with the `protected_content` keyword. For more information, see Using the protected_content Keyword, page 30-15.

Use the **Hash Type** drop-down to identify the hash function you used to encode your search string. The system supports SHA-512, SHA-256, and MD5 hashing for `protected_content` search strings. If the length of your hashed content does not match the selected hash type, the system does **not** save the rule.

The system automatically selects the Cisco-set default value. When **Default** is selected, no specific hash function is written into the rule and the system assumes SHA-512 for the hash function.

To specify a hash function when doing a protected content search:

Step 1 From the **Hash Type** drop-down list, select **Default**, **SHA-512**, **SHA-256**, or **MD5** as the hash for the `protected_content` keyword you are adding.

Tip If you select the Cisco-set **Default**, the system assumes SHA-512 as the hash function. For more information, see Hash Type, page 30-18.

Step 2 Continue with creating or editing the rule. See Constraining Content Matches, Searching for Content Matches, page 30-15, Writing New Rules, page 30-101, or Modifying Existing Rules, page 30-103 for more information.

Raw Data

License: Protection

The **Raw Data** option instructs the rules engine to analyze the original packet payload before analyzing the normalized payload data (decoded by a network analysis policy) and does not use an argument value. You can use this keyword when analyzing telnet traffic to check the telnet negotiation options in the payload before normalization.

You cannot use the **Raw Data** option together in the same `content` or `protected_content` keyword with any HTTP content option. See HTTP Content Options, page 30-23 for more information.

Tip You can configure the HTTP Inspect preprocessor **Client Flow Depth** and **Server Flow Depth** options to determine whether raw data is inspected in HTTP traffic, and how much raw data is inspected. For more information, see Selecting Server-Level HTTP Normalization Options, page 22-33.

To analyze raw data:

Step 1 Select the **Raw Data** check box for the `content` or `protected_content` keyword you are adding.

Step 2 Continue with creating or editing the rule. See Constraining Content Matches, Searching for Content Matches, page 30-15, Writing New Rules, page 30-101, or Modifying Existing Rules, page 30-103 for more information.

Not

License: Protection

Select the **Not** option to search for content that does not match the specified content. If you create a rule that includes a `content` or `protected_content` keyword with the **Not** option selected, you must also include in the rule at least one other `content` or `protected_content` keyword without the **Not** option selected.

Caution Do not create a rule that includes only one `content` or `protected_content` keyword if that keyword has the **Not** option selected. You may invalidate your intrusion policy.

For example, SMTP rule 1:2541:9 includes three `content` keywords, one of which has the **Not** option selected. A custom rule based on this rule would be invalid if you removed all of the `content` keywords except the one with the **Not** option selected. Adding such a rule to your intrusion policy could invalidate the policy.

To search for content that does not match the specified content:

Step 1 Select the **Not** check box for the `content` or `protected_content` keyword you are adding.

Tip You cannot select the **Not** check box and the **Use Fast Pattern Matcher** check box with the same `content` keyword.

Step 2 Include in the rule at least one other `content` or `protected_content` keyword that does not have the **Not** option selected.

Step 3 Continue with creating or editing the rule. See Constraining Content Matches, Searching for Content Matches, page 30-15, Writing New Rules, page 30-101, or Modifying Existing Rules, page 30-103 for more information.

Search Location Options

License: Protection

You can use search location options to specify where to begin searching for the specified content and how far to continue searching. For details about each option, see:

- Depth, page 30-20
- Distance, page 30-20
- Length, page 30-21
- Offset, page 30-21
- Within, page 30-21

For information about how to use search location options within the `content` or `protected_content` keyword, see:

- Using Search Location Options in the content Keyword, page 30-21
- Using Search Location Options in the protected_content Keyword, page 30-22

Depth



Note This option is **only** supported when configuring the `content` keyword. For more information, see Using the content Keyword, page 30-15.

Specifies the maximum content search depth, in bytes, from the beginning of the offset value, or if no offset is configured, from the beginning of the packet payload.

For example, in a rule with a content value of `cgi-bin/phf`, and `offset` value of `3`, and a `depth` value of `22`, the rule starts searching for a match to the `cgi-bin/phf` string at byte 3, and stops after processing 22 bytes (byte 25) in packets that meet the parameters specified by the rule header.

You must specify a value that is greater than or equal to the length of the specified content, up to a maximum of 65535 bytes. You cannot specify a value of 0.

The default depth is to search to the end of the packet.

Distance

Instructs the rules engine to identify subsequent content matches that occur a specified number of bytes after the previous successful content match.

Because the distance counter starts at byte 0, specify one less than the number of bytes you want to move forward from the last successful content match. For example, if you specify 4, the search begins at the fifth byte.

You can specify a value of -65535 to 65535 bytes. If you specify a negative `Distance` value, the byte you start searching on may fall outside the beginning of a packet. Any calculations will take into account the bytes outside the packet, even though the search actually starts on the first byte in the packet. For example, if the current location in the packet is the fifth byte, and the next content rule option specifies a `Distance` value of -10 and a `Within` value of 20, the search starts at the beginning of the payload and the `Within` option is adjusted to 15.

The default distance is 0, meaning the current location in the packet subsequent to the last content match.

Length



Note This option is **only** supported when configuring the `protected_content` keyword. For more information, see Using the protected_content Keyword, page 30-15.

The **Length** `protected_content` keyword option indicates the length, in bytes, of the unhashed search string.

For example, if you used the content `Sample1` to generate a secure hash, use `7` for the **Length** value. You **must** enter a value in this field.

Offset

Specifies in bytes where in the packet payload to start searching for content relative to the beginning of the packet payload. You can specify a value of-65535 to 65535 bytes.

Because the offset counter starts at byte 0, specify one less than the number of bytes you want to move forward from the beginning of the packet payload. For example, if you specify 7, the search begins at the eighth byte.

The default offset is 0, meaning the beginning of the packet.

Within



Note This option is **only** supported when configuring the `content` keyword. For more information, see Using the content Keyword, page 30-15.

The **Within** option indicates that, to trigger the rule, the next content match must occur within the specified number of bytes after the end of the last successful content match. For example, if you specify a **Within** value of `8`, the next content match must occur within the next eight bytes of the packet payload or it does not meet the criteria that triggers the rule.

You can specify a value that is greater than or equal to the length of the specified content, up to a maximum of 65535 bytes.

The default for **Within** is to search to the end of the packet.

Using Search Location Options in the content Keyword

You can use either of two `content` location pairs to specify where to begin searching for the specified content and how far to continue searching, as follows:

- Use **Offset** and **Depth** together to search relative to the beginning of the packet payload.
- Use **Distance** and **Within** together to search relative to the current search location.

When you specify only one of a pair, the default for the other option in the pair is assumed.

You cannot mix the **Offset** and **Depth** options with the **Distance** and **Within** options. For example, you cannot pair **Offset** and **Within**. You can use any number of location options in a rule.

When no location is specified, the defaults for **Offset** and **Depth** are assumed; that is, the content search starts at the beginning of the packet payload and continues to the end of the packet.

You can also use an existing `byte_extract` variable to specify the value for a location option. See Reading Packet Data into Keyword Arguments, page 30-82 for more information.

To specify a search location value in a content keyword:

Step 1 Type the value in the field for the `content` keyword you are adding. You have the following choices:

- **Offset**
- **Depth**
- **Distance**
- **Within**

You can use any number of location options in a rule.

Step 2 Continue with creating or editing the rule. See Constraining Content Matches, page 30-17, Searching for Content Matches, page 30-15, Writing New Rules, page 30-101 or Modifying Existing Rules, page 30-103 for more information.

Using Search Location Options in the protected_content Keyword

Use the required **Length** `protected_content` location option in combination with either the **Offset** or **Distance** location option to specify where to begin searching for the specified content and how far to continue searching, as follows:

- Use **Length** and **Offset** together to search for the protected string relative to the beginning of the packet payload.
- Use **Length** and **Distance** together to search for the protected string relative to the current search location.



Tip You cannot mix the **Offset** and **Distance** options within a single keyword configuration, but you can use any number of location options in a rule.

When no location is specified, the defaults are assumed; that is, the content search starts at the beginning of the packet payload and continues to the end of the packet.

You can also use an existing `byte_extract` variable to specify the value for a location option. For more information, see Reading Packet Data into Keyword Arguments, page 30-82.

To specify a search location value in a protected_content keyword:

Step 1 Type the value in the field for the `protected_content` keyword you are adding. You have the following choices:

- **Length** (required)
- **Offset**

- **Distance**

You cannot mix the **Offset** and **Distance** options within a single `protected_content` keyword, but you can use any number of location options in a rule.

Step 2 Continue with creating or editing the rule. See Constraining Content Matches, page 30-17, Searching for Content Matches, page 30-15, Writing New Rules, page 30-101 or Modifying Existing Rules, page 30-103 for more information.

HTTP Content Options

License: Protection

HTTP `content` or `protected_content` keyword options let you specify where to search for content matches within an HTTP message decoded by the HTTP Inspect preprocessor.

Two options search status fields in HTTP responses:

- **HTTP Status Code**
- **HTTP Status Message**

Note that although the rules engine searches the raw, unnormalized status fields, these options are listed here separately to simplify explanation below of the restrictions to consider when combining other raw HTTP fields and normalized HTTP fields.

Five options search normalized fields in HTTP requests, responses, or both, as appropriate (see HTTP Content Options, page 30-23 for more information):

- **HTTP URI**
- **HTTP Method**
- **HTTP Header**
- **HTTP Cookie**
- **HTTP Client Body**

Three options search raw (unnormalized) non-status fields in HTTP requests, responses, or both, as appropriate (see HTTP Content Options, page 30-23 for more information):

- **HTTP Raw URI**
- **HTTP Raw Header**
- **HTTP Raw Cookie**

Use the following guidelines when selecting HTTP `content` options:

- HTTP `content` options apply only to TCP traffic.
- To avoid a negative impact on performance, select only those parts of the message where the specified content might appear.

  For example, when traffic is likely to include large cookies such as those in shopping cart messages, you might search for the specified content in the HTTP header but not in HTTP cookies.
- To take advantage of HTTP Inspect preprocessor normalization, and to improve performance, any HTTP-related rule you create should at a minimum include at least one `content` or `protected_content` keyword with an **HTTP URI**, **HTTP Method**, **HTTP Header**, or **HTTP Client Body** option selected.
- You cannot use the `replace` keyword in conjunction with HTTP `content` or `protected_content` keyword options.

You can specify a single normalized HTTP option or status field, or use normalized HTTP options and status fields in any combination to target a content area to match. However, note the following restrictions when using HTTP field options:

- You cannot use the **Raw Data** option together in the same `content` or `protected_content` keyword with any HTTP option.
- You cannot use a raw HTTP field option (**HTTP Raw URI**, **HTTP Raw Header**, or **HTTP Raw Cookie**) together in the same `content` or `protected_content` keyword with its normalized counterpart (**HTTP URI**, **HTTP Header**, or **HTTP Cookie**, respectively).
- You cannot select **Use Fast Pattern Matcher** in combination with one or more of the following HTTP field options:

  HTTP Raw URI, **HTTP Raw Header**, **HTTP Raw Cookie**, **HTTP Cookie**, **HTTP Method**, **HTTP Status Message**, or **HTTP Status Code**

  However, you can include the options above in a `content` or `protected_content` keyword that also uses the fast pattern matcher to search one of the following normalized fields:

  HTTP URI, **HTTP Header**, or **HTTP Client Body**

  For example, if you select **HTTP Cookie**, **HTTP Header**, and **Use Fast Pattern Matcher**, the rules engine searches for content in both the HTTP cookie and the HTTP header, but the fast pattern matcher is applied only to the HTTP header, not to the HTTP cookie.
- When you combine restricted and unrestricted options, the fast pattern matcher searches only the unrestricted fields you specify to test whether to pass the rule to the rule editor for complete evaluation, including evaluation of the restricted fields. See Use Fast Pattern Matcher, page 30-26 for more information.

The above restrictions are reflected in the description of each option in the following list describing the HTTP `content` and `protected_content` keyword options.

HTTP URI

Select this option to search for content matches in the normalized request URI field.

Note that you cannot use this option in combination with the `pcre` keyword HTTP URI (U) option to search the same content. See the Snort-Specific Post Regular Expression Modifiers table for more information.

Note A pipelined HTTP request packet contains multiple URIs. When **HTTP URI** is selected and the rules engine detects a pipelined HTTP request packet, the rules engine searches all URIs in the packet for a content match.

HTTP Raw URI

Select this option to search for content matches in the normalized request URI field.

Note that you cannot use this option in combination with the `pcre` keyword HTTP URI (U) option to search the same content. See the Snort-Specific Post Regular Expression Modifiers table for more information.

Note A pipelined HTTP request packet contains multiple URIs. When **HTTP URI** is selected and the rules engine detects a pipelined HTTP request packet, the rules engine searches all URIs in the packet for a content match.

HTTP Method

Select this option to search for content matches in the request method field, which identifies the action such as GET and POST to take on the resource identified in the URI.

HTTP Header

Select this option to search for content matches in the normalized header field, except for cookies, in HTTP requests; also in responses when the HTTP Inspect preprocessor **Inspect HTTP Responses** option is enabled.

Note that you cannot use this option in combination with the `pcre` keyword HTTP header (H) option to search the same content. See the Snort-Specific Post Regular Expression Modifiers table for more information.

HTTP Raw Header

Select this option to search for content matches in the raw header field, except for cookies, in HTTP requests; also in responses when the HTTP Inspect preprocessor I**nspect HTTP Responses** option is enabled.

Note that you cannot use this option in combination with the `pcre` keyword HTTP raw header (D) option to search the same content. See the Snort-Specific Post Regular Expression Modifiers table for more information.

HTTP Cookie

Select this option to search for content matches in any cookie identified in a normalized HTTP client request header; also in response set-cookie data when the HTTP Inspect preprocessor **Inspect HTTP Responses** option is enabled. Note that the system treats cookies included in the message body as body content.

You must enable the HTTP Inspect preprocessor **Inspect HTTP Cookies** option to search only the cookie for a match; otherwise, the rules engine searches the entire header, including the cookie. See Selecting Server-Level HTTP Normalization Options, page 22-33 for more information.

Note the following:

– You cannot use this option in combination with the `pcre` keyword HTTP cookie (C) option to search the same content. See the Snort-Specific Post Regular Expression Modifiers table for more information.

– The `Cookie:` and `Set-Cookie:` header names, leading spaces on the header line, and the `CRLF` that terminates the header line are inspected as part of the header and not as part of the cookie.

HTTP Raw Cookie

Select this option to search for content matches in any cookie identified in a raw HTTP client request header; also in response set-cookie data when the HTTP Inspect preprocessor **Inspect HTTP Responses** option is enabled; note that the system treats cookies included in the message body as body content.

You must enable the HTTP Inspect preprocessor **Inspect HTTP Cookies** option to search only the cookie for a match; otherwise, the rules engine searches the entire header, including the cookie. See Selecting Server-Level HTTP Normalization Options, page 22-33 for more information.

Note the following:

– You cannot use this option in combination with the `pcre` keyword HTTP raw cookie (K) option to search the same content. See the Snort-Specific Post Regular Expression Modifiers table for more information.

– The `Cookie:` and `Set-Cookie:` header names, leading spaces on the header line, and the `CRLF` that terminates the header line are inspected as part of the header and not as part of the cookie.

HTTP Client Body

Select this option to search for content matches in the message body in an HTTP client request.

Note that for this option to function, you must specify a value of 0 to 65535 for the HTTP Inspect preprocessor **HTTP Client Body Extraction Depth** option. See Selecting Server-Level HTTP Normalization Options, page 22-33 for more information.

HTTP Status Code

Select this option to search for content matches in the 3-digit status code in an HTTP response.

You must enable the HTTP Inspect preprocessor **Inspect HTTP Responses** option for this option to return a match. See Selecting Server-Level HTTP Normalization Options, page 22-33 for more information.

HTTP Status Message

Select this option to search for content matches in the textual description that accompanies the status code in an HTTP response.

You must enable the HTTP Inspect preprocessor **Inspect HTTP Responses** option for this option to return a match. See Selecting Server-Level HTTP Normalization Options, page 22-33 for more information.

To specify an HTTP content option when doing a content search of TCP traffic:

Step 1 Optionally, to take advantage of HTTP Inspect preprocessor normalization, and to improve performance, select:

- at least one from among the **HTTP URI**, **HTTP Raw URI**, **HTTP Method**, **HTTP Header**, **HTTP Raw Header**, or **HTTP Client Body** options for the `content` or `protected_content` keyword you are adding
- the **HTTP Cookie** or **HTTP Raw Cookie** option

Step 2 Continue with creating or editing the rule. See Constraining Content Matches, page 30-17, Searching for Content Matches, page 30-15, Writing New Rules, page 30-101, or Modifying Existing Rules, page 30-103 for more information.

Use Fast Pattern Matcher

License: Protection

Note These options are **not** supported when configuring the `protected_content` keyword. For more information, see Using the protected_content Keyword, page 30-15.

The fast pattern matcher quickly determines which rules to evaluate before passing a packet to the rules engine. This initial determination improves performance by significantly reducing the number of rules used in packet evaluation.

By default, the fast pattern matcher searches packets for the longest content specified in a rule; this is to eliminate as much as possible needless evaluation of a rule. Consider the following example rule fragment:

```
alert tcp any any -> any 80 (msg:"Exploit"; content:"GET";
http_method; nocase; content:"/exploit.cgi"; http_uri;
nocase;)
```

Almost all HTTP client requests contain the content `GET`, but few will contain the content `/exploit.cgi`. Using `GET` as the fast pattern content would cause the rules engine to evaluate this rule in most cases and would rarely result in a match. However, most client `GET` requests would not be evaluated using `/exploit.cgi`, thus increasing performance.

The rules engine evaluates the packet against the rule only when the fast pattern matcher detects the specified content. For example, if one `content` keyword in a rule specifies the content `short`, another specifies `longer`, and a third specifies `longest`, the fast pattern matcher will use the content `longest` and the rule will be evaluated only if the rules engine finds `longest` in the payload.

You can use the **Use Fast Pattern Matcher** option to specify a shorter search pattern for the fast pattern matcher to use. Ideally, the pattern you specify is less likely to be found in the packet than the longest pattern and, therefore, more specifically identifies the targeted exploit.

Note the following restrictions when selecting **Use Fast Pattern Matcher** and other options in the same `content` keyword:

- You can specify **Use Fast Pattern Matcher** only one time per rule.
- You cannot use **Distance**, **Within**, **Offset**, or **Depth** when you select **Use Fast Pattern Matcher** in combination with **Not**.
- You cannot select Use Fast Pattern Matcher in combination with any of the following HTTP field options:

  HTTP Raw URI, **HTTP Raw Header**, **HTTP Raw Cookie**, **HTTP Cookie**, **HTTP Method**, **HTTP Status Message**, or **HTTP Status Code**

  However, you can include the options above in a `content` keyword that also uses the fast pattern matcher to search one of the following normalized fields:

  HTTP URI, **HTTP Header**, or **HTTP Client Body**

  For example, if you select **HTTP Cookie**, **HTTP Header**, and **Use Fast Pattern Matcher**, the rules engine searches for content in both the HTTP cookie and the HTTP header, but the fast pattern matcher is applied only to the HTTP header, not to the HTTP cookie.

  Note that you cannot use a raw HTTP field option (**HTTP Raw URI**, **HTTP Raw Header**, or **HTTP Raw Cookie**) together in the same `content` keyword with its normalized counterpart (**HTTP URI**, **HTTP Header**, or **HTTP Cookie**, respectively). See HTTP Content Options, page 30-23 for more information.

  When you combine restricted and unrestricted options, the fast pattern matcher searches only the unrestricted fields you specify to test whether to pass the packet to the rules engine for complete evaluation, including evaluation of the restricted fields.
- Optionally, when you select **Use Fast Pattern Matcher** you can also select **Fast Pattern Matcher Only** or **Fast Pattern Matcher Offset and Length**, but not both.
- You cannot use the fast pattern matcher when inspecting Base64 data; see Decoding and Inspecting Base64 Data, page 30-99 for more information.

Using the Fast Pattern Matcher Only

The **Fast Pattern Matcher Only** option allows you to use the `content` keyword only as a fast pattern matcher option and not as a rule option. You can use this option to conserve resources when rules engine evaluation of the specified content is not necessary. For example, consider a case where a rule requires only that the content `12345` be anywhere in the payload. When the fast pattern matcher detects the pattern, the packet can be evaluated against additional keywords in the rule. There is no need for the rules engine to reevaluate the packet to determine if it includes the pattern `12345`.

You would not use this option when the rule contains other conditions relative to the specified content. For example, you would not use this option to search for the content `1234` if another rule condition sought to determine if `abcd` occurs before `1234`. In this case, the rules engine could not determine the relative location because specifying **Fast Pattern Matcher Only** instructs the rules engine not to search for the specified content.

Note the following conditions when using this option:

- The specified content is location-independent; that is, it may occur anywhere in the payload; thus, you cannot use positional options (**Distance**, **Within**, **Offset**, **Depth**, or **Fast Pattern Matcher Offset and Length**).
- You cannot use this option in combination with **Not**.
- You cannot use this option in combination with **Fast Pattern Matcher Offset and Length**.
- The specified content will be treated as case-insensitive, because all patterns are inserted into the fast pattern matcher in a case-insensitive manner; this is handled automatically, so it is not necessary to select **Case Insensitive** when you select this option.
- You should not immediately follow a `content` keyword that uses the **Fast Pattern Matcher Only** option with the following keywords, which set the search location relative to the current search location:
- `isdataat`
- `pcre`
- `content` when **Distance** or **Within** is selected
- `content` when **HTTP URI** is selected
- `asn1`
- `byte_jump`
- `byte_test`
- `byte_extract`
- `base64_decode`

Specifying Fast Pattern Matcher Offset and Length

The **Fast Pattern Matcher Offset and Length** option allows you to specify a portion of the content to search. This can reduce memory consumption in cases where the pattern is very long and only a portion of the pattern is sufficient to identify the rule as a likely match. When a rule is selected by the fast pattern matcher, the entire pattern is evaluated against the rule.

You determine the portion for the fast pattern matcher to use by specifying in bytes where to begin the search (offset) and how far into the content (length) to search, using the syntax:

```
offset,length
```

For example, for the content:

```
1234567
```

if you specify the number of offset and length bytes as:

```
1,5
```

the fast pattern matcher searches only for the content `23456`.

Note that you cannot use this option together with **Fast Pattern Matcher Only**.

To specify the content searched for by the fast pattern matcher:

Step 1 Select **Use Fast Pattern Matcher** for the `content` keyword you are adding.

Step 2 Optionally, select **Fast Pattern Matcher Only** to determine without rules engine evaluation if the specified pattern exists in the packet.

Evaluation proceeds only if the fast pattern matcher detects the specified content.

Step 3 Optionally, specify in **Fast Pattern Matcher Offset and Length** a portion of the pattern to search for the content using the syntax:

```
offset,length
```

where *`offset`* specifies how many bytes from the beginning of the content to begin the search, and *`length`* specifies the number of bytes to continue.

Step 4 Continue with creating or editing the rule. See Constraining Content Matches, page 30-17, Searching for Content Using PCRE, page 30-35, Writing New Rules, page 30-101, or Modifying Existing Rules, page 30-103 for more information.

Replacing Content in Inline Deployments

License: Protection

You can use the `replace` keyword in an inline deployment to replace specified content.

Note You **cannot** use the `replace` keyword to replace content in SSL traffic detected by the Cisco SSL Appliance. The original encrypted data, not the replacement data, will be transmitted. See the *Cisco SSL Appliance Administration and Deployment Guide* for more information.

To use the `replace` keyword, construct a custom standard text rule that uses the `content` keyword to look for a specific string. Then use the `replace` keyword to specify a string to replace the content. The replace value and content value must be the same length.

Note You **cannot** use the `replace` keyword to replace hashed content in a `protected_content` keyword. For more information, see Using the protected_content Keyword, page 30-15.

Optionally, you can enclose the replacement string in quotation marks for backward compatibility with previous ASA FirePOWER module software versions. If you do not include quotation marks, they are added to the rule automatically so the rule is syntactically correct. To include a leading or trailing quotation mark as part of the replacement text, you must use a backslash to escape it, as shown in the following example:

```
"replacement text plus \"quotation\" marks""
```

A rule can contain multiple `replace` keywords, but only one per `content` keyword. Only the first instance of the content found by the rule is replaced.

The following explain example uses of the `replace` keyword:

- If the system detects an incoming packet that contains an exploit, you can replace the malicious string with a harmless one. Sometimes this technique is more successful than simply dropping the offending packet. In some attack scenarios, the attacker simply resends the dropped packet until it bypasses your network defenses or floods your network. By substituting one string for another rather than dropping the packet, you may trick the attacker into believing that the attack was launched against a target that was not vulnerable.
- If you are concerned about reconnaissance attacks that try to learn whether you are running a vulnerable version of, for example, a web server, then you can detect the outgoing packet and replace the banner with your own text.



Note Make sure that you set the rule state to Generate Events in the inline intrusion policy where you want to use the replace rule; setting the rule to Drop and Generate events would cause the packet to drop, which would prevent replacing the content.

As part of the string replacement process, the system automatically updates the packet checksums so that the destination host can receive the packet without error.

Note that you cannot use the `replace` keyword in combination with HTTP request message `content` keyword options. See Searching for Content Matches, page 30-15 and HTTP Content Options, page 30-23 for more information.

To replace content in an inline deployment:

Step 1 On the Create Rule page, select **content** in the drop-down list and click **Add Option.**

The `content` keyword appears.

Step 2 Specify the content you want to detect in the **content** field and, optionally, select any applicable arguments. Note that you cannot use the HTTP request message `content` keyword options with the `replace` keyword.

Step 3 Select **replace** in the drop-down list and click **Add Option.**

The `replace` keyword appears beneath the `content` keyword.

Step 4 Specify the replacement string for the specified content in the **replace:** field.

Using Byte_Jump and Byte_Test

License: Protection

You can use `byte_jump` and `byte_test` to calculate where in a packet the rules engine should begin testing for a data match, and which bytes it should evaluate.

You can also use the `byte_jump` and `byte_test` **DCE/RPC** argument to tailor either keyword for traffic processed by the DCE/RPC preprocessor. When you use the **DCE/RPC** argument, you can also use `byte_jump` and `byte_test` in conjunction with other specific DCE/RPC keywords. See Decoding DCE/RPC Traffic, page 22-2 and DCE/RPC Keywords, page 30-59 for more information.

See the following sections for more information:

- byte_jump, page 30-31
- byte_test, page 30-33

byte_jump

License: Protection

The `byte_jump` keyword calculates the number of bytes defined in a specified byte segment, and then skips that number of bytes within the packet, either forward from the end of the specified byte segment, or from the beginning of the packet payload, depending on the options you specify. This is useful in packets where a specific segment of bytes describe the number of bytes included in variable data within the packet.

The following table describes the arguments required by the `byte_jump` keyword.

Table 30-8 Required byte_jump Arguments

Argument	Description
Bytes	The number of bytes to calculate from the packet.
Offset	The number of bytes into the payload to start processing. The `offset` counter starts at byte 0, so calculate the `offset` value by subtracting 1 from the number of bytes you want to jump forward from the beginning of the packet payload or the last successful content match. You can also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments, page 30-82 for more information.

The following table describes options you can use to define how the system interprets the values you specified for the required arguments.

Table 30-9 Additional Optional byte_jump Arguments

Argument	Description
Relative	Makes the offset relative to the last pattern found in the last successful content match.
Align	Rounds the number of converted bytes up to the next 32-bit boundary.
Multiplier	Indicates the value by which the rules engine should multiply the `byte_jump` value obtained from the packet to get the final `byte_jump` value. That is, instead of skipping the number of bytes defined in a specified byte segment, the rules engine skips that number of bytes multiplied by an integer you specify with the Multiplier argument.
Post Jump Offset	The number of bytes -63535 through 63535 to skip forward or backward after applying other `byte_jump` arguments. A positive value skips forward and a negative value skips backward. Leave the field blank or enter `0` to disable. See the **DCE/RPC** argument in the Endianness Arguments table for `byte_jump` arguments that do not apply when you select the **DCE/RPC** argument.
From Beginning	Indicates that the rules engine should skip the specified number of bytes in the payload starting from the beginning of the packet payload, rather than from the end of the byte segment that specifies the number of bytes to skip.

You can specify only one of **DCE/RPC**, **Endian**, or **Number Type**.

If you want to define how the `byte_jump` keyword calculates the bytes, you can choose from the arguments described in the following table (if neither argument is specified, network byte order is used).

Table 30-10 Endianness Arguments

Argument	Description
Big Endian	Processes data in big endian byte order, which is the default network byte order.
Little Endian	Processes data in little endian byte order.
DCE/RPC	Specifies a `byte_jump` keyword for traffic processed by the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic, page 22-2 for more information. The DCE/RPC preprocessor determines big endian or little endian byte order, and the **Number Type**, **Endian**, and **From Beginning** arguments do not apply. When you enable this argument, you can also use `byte_jump` in conjunction with other specific DCE/RPC keywords. See DCE/RPC Keywords, page 30-59 for more information.

Define how the system views string data in a packet by using one of the arguments in the following table.

Table 30-11 Number Type Arguments

Argument	Description
Hexadecimal String	Represents converted string data in hexadecimal format.
Decimal String	Represents converted string data in decimal format.
Octal String	Represents converted string data in octal format.

For example, if the values you set for `byte_jump` are as follows:

- Bytes = 4
- Offset = 12
- Relative enabled
- Align enabled

the rules engine calculates the number described in the four bytes that appear 13 bytes after the last successful content match, and skips ahead that number of bytes in the packet. For instance, if the four calculated bytes in a specific packet were `00 00 00 1F`, the rules engine would convert this to 31. Because `align` is specified (which instructs the engine to move to the next 32-bit boundary), the rules engine skips ahead 32 bytes in the packet.

Alternately, if the values you set for `byte_jump` are as follows:

- Bytes = 4
- Offset = 12
- From Beginning enabled
- Multiplier = 2

the rules engine calculates the number described in the four bytes that appear 13 bytes after the beginning of the packet. Then, the engine multiplies that number by two to obtain the total number of bytes to skip. For instance, if the four calculated bytes in a specific packet were `00 00 00 1F`, the rules engine would convert this to 31, then multiply it by two to get 62. Because From Beginning is enabled, the rules engine skips the first 63 bytes in the packet.

To use byte_jump:

Step 1 Select `byte_jump` in the drop-down list and click **Add Option**.

The byte_jump section appears beneath the last keyword you selected.

byte_test

License: Protection

The `byte_test` keyword calculates the number of bytes in a specified byte segment and compares them, according to the operator and value you specify.

The following table describes the required arguments for the `byte_test` keyword.

Table 30-12 Required byte_test Arguments

Argument	Description
Bytes	The number of bytes to calculate from the packet. You can specify 1 to 10 bytes.
Operator and Value	Compares the specified value to <, >, =, !, &, ^, !>, !<, !=, !&, or !^. For example, if you specify `!1024`, `byte_test` would convert the specified number, and if it did not equal 1024, it would generate an event (if all other keyword parameters matched). Note that ! and != are equivalent. You can also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments, page 30-82 for more information.
Offset	The number of bytes into the payload to start processing. The `offset` counter starts at byte 0, so calculate the `offset` value by subtracting 1 from the number of bytes you want to count forward from the beginning of the packet payload or the last successful content match. You can also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments, page 30-82 for more information.

You can further define how the system uses `byte_test` arguments with the arguments described in the following table.

Table 30-13 Additional Optional byte_test Arguments

Argument	Description
Relative	Makes the offset relative to the last successful pattern match.
Align	Rounds the number of converted bytes up to the next 32-bit boundary.

You can specify only one of **DCE/RPC**, **Endian**, or **Number Type**.

To define how the `byte_test` keyword calculates the bytes it tests, choose from the arguments in the following table. If neither argument is specified, network byte order is used.

Table 30-14 Endianness byte_test Arguments

Argument	Description
Big Endian	Processes data in big endian byte order, which is the default network byte order.
Little Endian	Processes data in little endian byte order.
DCE/RPC	Specifies a `byte_test` keyword for traffic processed by the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic, page 22-2 for more information. The DCE/RPC preprocessor determines big endian or little endian byte order, and the **Number Type** and **Endian** argument do not apply. When you enable this argument, you can also use `byte_test` in conjunction with other specific DCE/RPC keywords. See DCE/RPC Keywords, page 30-59 for more information.

You can define how the system views string data in a packet by using one of the arguments in the following table.

Table 30-15 Number Type byte-test Arguments

Argument	Description
Hexadecimal String	Represents converted string data in hexadecimal format.
Decimal String	Represents converted string data in decimal format.
Octal String	Represents converted string data in octal format.

For example, if the value for `byte_test` is specified as the following:

- Bytes = 4
- Operator and Value > 128
- Offset = 8
- Relative enabled

the rules engine calculates the number described in the four bytes that appear 9 bytes away from (relative to) the last successful content match, and, if the calculated number is larger than 128 bytes, the rule is triggered.

To use byte_test:

Step 1 On the Create Rule page, select `byte_test` in the drop-down list and click **Add Option**.

The `byte_test` section appears beneath the last keyword you selected.

Searching for Content Using PCRE

License: Protection

The `pcre` keyword allows you to use Perl-compatible regular expressions (PCRE) to inspect packet payloads for specified content. You can use PCRE to avoid writing multiple rules to match slight variations of the same content.

Regular expressions are useful when searching for content that could be displayed in a variety of ways. The content may have different attributes that you want to account for in your attempt to locate it within a packet's payload.

Note that the regular expression syntax used in intrusion rules is a subset of the full regular expression library and varies in some ways from the syntax used in commands in the full library. When adding a `pcre` keyword using the rule editor, enter the full value in the following format:

`!/pcre/ ismxAEGRBUIPHDMCKSY`

where:

- ! is an optional negation (use this if you want to match patterns that **do not** match the regular expression).
- `/pcre/` is a Perl-compatible regular expression.
- `ismxAEGRBUIPHDMCKSY` is any combination of modifier options.

Also note that you must escape the characters listed in the following table for the rules engine to interpret them correctly when you use them in a PCRE to search for specific content in a packet payload.

Table 30-16 Escaped PCRE Characters

You must escape...	with a backslash...	or Hex code...
# (hash mark)	\#	\x23
; (semicolon)	\;	\x3B
| (vertical bar)	|	\x7C
: (colon)	\:	\x3A

Tip Optionally, you can surround your Perl-compatible regular expression with quote characters, for example, *`pcre_expression`* or *`"pcre_expression"`*. The option of using quotes accommodates experienced users accustomed to previous versions when quotes were required instead of optional. The rule editor does not display quotation marks when you display a rule after saving it.

You can also use `m?regex?`, where `?` is a delimiter other than /. You may want to use this in situations where you need to match a forward slash within a regular expression and do not want to escape it with a backslash. For example, you might use `m?`*`regex`*`? ismxAEGRBUIPHDMCKSY` where *`regex`* is your

Perl-compatible regular expression and `ismxAEGRBUIPHDMCKSY` is any combination of modifier options. See Perl-Compatible Regular Expression Basics, page 30-36 for more information about regular expression syntax.

The following sections provide more information about building valid values for the `pcre` keyword:

- Perl-Compatible Regular Expression Basics, page 30-36 describes the common syntax used in Perl-compatible regular expressions.
- PCRE Modifier Options, page 30-37 describes the options you can use to modify your regular expression.
- Example PCRE Keyword Values, page 30-41 gives example usage of the `pcre` keyword in rules.

Perl-Compatible Regular Expression Basics

License: Protection

The `pcre` keyword accepts standard Perl-compatible regular expression (PCRE) syntax. The following sections describe that syntax.

Tip While this section describes the basic syntax you may use for PCRE, you may want to consult an online reference or book dedicated to Perl and PCRE for more advanced information.

Metacharacters

License: Protection

Metacharacters are literal characters that have special meaning within regular expressions. When you use them within a regular expression, you must "escape" them by preceding them with a backslash.

The following table describes the metacharacters you can use with PCRE and gives examples of each.

Table 30-17 PCRE Metacharacters

Metacharacter	Description	Example
`.`	Matches any character except newlines. If `s` is used as a modifying option, it also includes newline characters.	`abc.` matches `abcd`, `abc1`, `abc#`, and so on.
`*`	Matches zero or more occurrences of a character or expression.	`abc*` matches `abc`, `abcc`, `abccc`, `abcccc`, and so on.
`?`	Matches zero or one occurrence of a character or expression.	`abc?` matches `abc`.
`+`	Matches one or more occurrences of a character or expression.	`abc+` matches `abc`, `abcc`, `abccc`, `abcccc`, and so on.
`()`	Groups expressions.	`(abc)+` matches `abc`, `abcabc`, `abcabcabc` and so on.
`{}`	Specifies a limit for the number of matches for a character or expression. If you want to set a lower and upper limit, separate the lower limit and upper limit with a comma.	`a{4,6}` matches `aaaa`, `aaaaa`, or `aaaaaa`. `(ab){2}` matches `abab`.
`[]`	Allows you to define character classes, and matches any character or combination of characters described in the set.	`[abc123]` matches `a` or `b` or `c`, and so on.

Table 30-17 PCRE Metacharacters (continued)

Metacharacter	Description	Example
^	Matches content at the beginning of a string. Also used for negation, if used within a character class.	^in matches the "in" in info, but not in bin. [^a] matches anything that does not contain a.
$	Matches content at the end of a string.	ce$ matches the "ce" in announce, but not cent.
|	Indicates an OR expression.	(MAILTO|HELP) matches MAILTO or HELP.
\	Allows you to use metacharacters as actual characters and is also used to specify a predefined character class.	\. matches a period, * matches an asterisk, \\ matches a backslash and so on. \d matches the numeric characters, \w matches alphanumeric characters, and so on. See Character Classes, page 30-37 for more information about using character classes in PCRE.

Character Classes

License: Protection

Character classes include alphabetic characters, numeric characters, alphanumeric characters, and white space characters. While you can create your own character classes within brackets (see Metacharacters, page 30-36), you can use the predefined classes as shortcuts for different types of character types. When used without additional qualifiers, a character class matches a single digit or character.

The following table describes and provides examples of the predefined character classes accepted by PCRE.

Table 30-18 PCRE Character Classes

Character Class	Description	Character Class Definition
\d	Matches a numeric character ("digit").	[0-9]
\D	Matches anything that is not an numeric character.	[^0-9]
\w	Matches an alphanumeric character ("word").	[a-zA-Z0-9_]
\W	Matches anything that is not an alphanumeric character.	[^a-zA-Z0-9_]
\s	Matches white space characters, including spaces, carriage returns, tabs, newlines, and form feeds.	[\r\t\n\f]
\S	Matches anything that is not a white space character.	[^ \r\t\n\f]

PCRE Modifier Options

License: Protection

You can use modifying options after you specify regular expression syntax in the pcre keyword's value. These modifiers perform Perl, PCRE, and Snort-specific processing functions. Modifiers always appear at the end of the PCRE value, and appear in the following format:

/*pcre*/ismxAEGRBUIPHDMCKSY

where ismxAEGRBUPHMC can include any of the modifying options that appear in the following tables.

Tip Optionally, you can surround the regular expression and any modifying options with quotes, for example, `"/pcre/ismxAEGRBUIPHDMCKSY"`. The option of using quotes accommodates experienced users accustomed to previous versions when quotes were required instead of optional. The rule editor does not display quotation marks when you display a rule after saving it.

The following table describes options you can use to perform Perl processing functions.

Table 30-19 Perl-Related Post Regular Expression Options

Option	Description
i	Makes the regular expression case-insensitive.
s	The dot character (.) describes all characters except the newline or `\n` character. You can use `"s"` as an option to override this and have the dot character match all characters, including the newline character.
m	By default, a string is treated as a single line of characters, and `^` and `$` match the beginning and ending of a specific string. When you use `"m"` as an option, `^` and `$` match content immediately before or after any newline character in the buffer, as well as at the beginning or end of the buffer.
x	Ignores white space data characters that may appear within the pattern, except when escaped (preceded by a backslash) or included inside a character class.

The following table describes the PCRE modifiers you can use after the regular expression.

Table 30-20 PCRE-Related Post Regular Expression Options

Option	Description
A	The pattern must match at the beginning of the string (same as using `^` in a regular expression).
E	Sets `$` to match only at the end of the subject string. (Without `E`, `$` also matches immediately before the final character if it is a newline, but not before any other newline characters).
G	By default, `* +` and `?` are "greedy," which means that if two or more matches are found, they will choose the longest match. Use the G character to change this so that these characters always choose the first match unless followed by a question mark character (`?`). For example, `*? +?` and `??` would be greedy in a construct using the G modifier, and any incidences of `*`, `+`, or `?` without the additional question mark will not be greedy.

The following table describes the Snort-specific modifiers that you can use after the regular expression.

Table 30-21 Snort-Specific Post Regular Expression Modifiers

Option	Description
R	Searches for matching content relative to the end of the last match found by the rules engine.
B	Searches for the content within data before it is decoded by a preprocessor (this option is similar to using the `Raw Data` argument with the `content` or `protected_content` keyword).

Table 30-21 Snort-Specific Post Regular Expression Modifiers (continued)

Option	Description
U	Searches for the content within the URI of a normalized HTTP request message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP URI** option to search the same content. See HTTP Content Options, page 30-23 for more information. **Note** A pipelined HTTP request packet contains multiple URIs. A PCRE expression that includes the U option causes the rules engine to search for a content match only in the first URI in a pipelined HTTP request packet. To search all URIs in the packet, use the `content` or `protected_content` keyword with **HTTP URI** selected, either with or without an accompanying PCRE expression that uses the U option.
I	Searches for the content within the URI of a raw HTTP request message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Raw URI** option to search the same content. See HTTP Content Options, page 30-23 for more information.
P	Searches for the content within the body of a normalized HTTP request message decoded by the HTTP Inspect preprocessor. See the `content` and `protected_content` keyword **HTTP Client Body** option in HTTP Content Options, page 30-23 for more information.
H	Searches for the content within the header, excluding cookies, of an HTTP request or response message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Header** option to search the same content. See HTTP Content Options, page 30-23 for more information.
D	Searches for the content within the header, excluding cookies, of a raw HTTP request or response message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Raw Header** option to search the same content. See HTTP Content Options, page 30-23 for more information.
M	Searches for the content within the method field of a normalized HTTP request message decoded by the HTTP Inspect preprocessor; the method field identifies the action such as GET, PUT, CONNECT, and so on to take on the resource identified in the URI. See the `content` and `protected_content` keyword **HTTP Method** option in HTTP Content Options, page 30-23 for more information.
C	When the HTTP Inspect preprocessor **Inspect HTTP Cookies** option is enabled, searches for the normalized content within any cookie in an HTTP request header, and also within any set-cookie in an HTTP response header when the preprocessor **Inspect HTTP Responses** option is enabled. When **Inspect HTTP Cookies** is not enabled, searches the entire header, including the cookie or set-cookie data. Note the following: • Cookies included in the message body are treated as body content. • You cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Cookie** option to search the same content. See HTTP Content Options, page 30-23 for more information. • The `Cookie:` and `Set-Cookie:` header names, leading spaces on the header line, and the `CRLF` that terminates the header line are inspected as part of the header and not as part of the cookie.

Table 30-21 Snort-Specific Post Regular Expression Modifiers (continued)

Option	Description
U	Searches for the content within the URI of a normalized HTTP request message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP URI** option to search the same content. See HTTP Content Options, page 30-23 for more information. **Note** A pipelined HTTP request packet contains multiple URIs. A PCRE expression that includes the U option causes the rules engine to search for a content match only in the first URI in a pipelined HTTP request packet. To search all URIs in the packet, use the `content` or `protected_content` keyword with **HTTP URI** selected, either with or without an accompanying PCRE expression that uses the U option.
I	Searches for the content within the URI of a raw HTTP request message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Raw URI** option to search the same content. See HTTP Content Options, page 30-23 for more information.
P	Searches for the content within the body of a normalized HTTP request message decoded by the HTTP Inspect preprocessor. See the `content` and `protected_content` keyword **HTTP Client Body** option in HTTP Content Options, page 30-23 for more information.
H	Searches for the content within the header, excluding cookies, of an HTTP request or response message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Header** option to search the same content. See HTTP Content Options, page 30-23 for more information.
D	Searches for the content within the header, excluding cookies, of a raw HTTP request or response message decoded by the HTTP Inspect preprocessor. Note that you cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Raw Header** option to search the same content. See HTTP Content Options, page 30-23 for more information.
M	Searches for the content within the method field of a normalized HTTP request message decoded by the HTTP Inspect preprocessor; the method field identifies the action such as GET, PUT, CONNECT, and so on to take on the resource identified in the URI. See the `content` and `protected_content` keyword **HTTP Method** option in HTTP Content Options, page 30-23 for more information.
C	When the HTTP Inspect preprocessor **Inspect HTTP Cookies** option is enabled, searches for the normalized content within any cookie in an HTTP request header, and also within any set-cookie in an HTTP response header when the preprocessor **Inspect HTTP Responses** option is enabled. When **Inspect HTTP Cookies** is not enabled, searches the entire header, including the cookie or set-cookie data. Note the following: • Cookies included in the message body are treated as body content. • You cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Cookie** option to search the same content. See HTTP Content Options, page 30-23 for more information. • The `Cookie:` and `Set-Cookie:` header names, leading spaces on the header line, and the `CRLF` that terminates the header line are inspected as part of the header and not as part of the cookie.

Table 30-21 Snort-Specific Post Regular Expression Modifiers (continued)

Option	Description
K	When the HTTP Inspect preprocessor **Inspect HTTP Cookies** option is enabled, searches for the raw content within any cookie in an HTTP request header, and also within any set-cookie in an HTTP response header when the preprocessor **Inspect HTTP Responses** option is enabled. When **Inspect HTTP Cookies** is not enabled, searches the entire header, including the cookie or set-cookie data. Note the following: • Cookies included in the message body are treated as body content. • You cannot use this option in combination with the `content` or `protected_content` keyword **HTTP Raw Cookie** option to search the same content. See HTTP Content Options, page 30-23 for more information. • The `Cookie:` and `Set-Cookie:` header names, leading spaces on the header line, and the `CRLF` that terminates the header line are inspected as part of the header and not as part of the cookie.
S	Searches the 3-digit status code in an HTTP response. See the `content` and `protected_content` keyword **HTTP Status Code** option in HTTP Content Options, page 30-23 for more information.
Y	Searches the textual description that accompanies the status code in an HTTP response. See the `content` and `protected_content` keyword **HTTP Status Message** option in HTTP Content Options, page 30-23 for more information.

Note Do not use the U option in combination with the R option. This could cause performance problems. Also, do not use the U option in combination with any other HTTP content option (I, P, H, D, M, C, K, S, or Y).

Example PCRE Keyword Values

License: Protection

The following examples show values that you could enter for `pcre`, with descriptions of what each example would match.

- **`/feedback[(\d{0,1})]?\.cgi/U`**

  This example searches packet payload for `feedback`, followed by zero or one numeric character, followed by `.cgi`, and located only in URI data.

  This example would match:

- `feedback.cgi`
- `feedback1.cgi`
- `feedback2.cgi`
- `feedback3.cgi`

  This example would **not** match:

- `feedbacka.cgi`
- `feedback11.cgi`
- `feedback21.cgi`
- `feedbackzb.cgi`
- **`/^ez(\w{3,5})\.cgi/iU`**

  This example searches packet payload for `ez` at the beginning of a string, followed by a word of 3 to 5 letters, followed by `.cgi`. The search is case-insensitive and only searches URI data.

This example would match:

- `EZBoard.cgi`
- `ezman.cgi`
- `ezadmin.cgi`
- `EZAdmin.cgi`

This example would **not** match:

- `ezez.cgi`
- `fez.cgi`
- `abcezboard.cgi`
- `ezboardman.cgi`
- **`/mail(file|seek)\.cgi/U`**

  This example searches packet payload for `mail`, followed by either `file` or `seek`, in URI data.

  This example would match:

- `mailfile.cgi`
- `mailseek.cgi`

  This example would **not** match:

- `MailFile.cgi`
- `mailfilefile.cgi`
- **`m?http\\x3a\x2f\x2f.*(\n|\t)+?U`**

  This example searches packet payload for URI content for a tab or newline character in an HTTP request, after any number of characters. This example uses `m?`*`regex`*`?` to avoid using `http\:\/\/` in the expression. Note that the colon is preceded by a backslash.

  This example would match:

- `http://www.example.com?scriptvar=x&othervar=\n\..\..`
- `http://www.example.com?scriptvar=\t`

  This example would **not** match:

- `ftp://ftp.example.com?scriptvar=&othervar=\n\..\..`
- `http://www.example.com?scriptvar=|/bin/sh -i|`
- **`m?http\\x3a\x2f\x2f.*=\|.*\|+?sU`**

  This example searches packet payload for a URL with any number of characters, including newlines, followed by an equal sign, and pipe characters that contain any number of characters or white space. This example uses `m?`*`regex`*`?` to avoid using `http\:\/\/` in the expression.

  This example would match:

- `http://www.example.com?value=|/bin/sh/ -i|`
- `http://www.example.com?input=|cat /etc/passwd|`

  This example would **not** match:

- `ftp://ftp.example.com?value=|/bin/sh/ -i|`
- `http://www.example.com?value=x&input?|cat /etc/passwd|`
- **`/[0-9a-f]{2}\:[0-9a-f]{2}\:[0-9a-f]{2}\:[0-9a-f]{2}\:[0-9a-f]{2}\:[0-9a-f]{2}/i`**

  This example searches packet payload for any MAC address. Note that it escapes the colon characters with backslashes.

Adding Metadata to a Rule

License: Protection

You can use the `metadata` keyword to add descriptive information to a rule. You can use the information you add to organize or identify rules in ways that suit your needs, and to search for rules.

The system validates metadata based on the format:

```
key value
```

where *key* and *value* provide a combined description separated by a space. This is the format used by the Cisco VRT for adding metadata to rules provided by Cisco.

Alternatively, you can also use the format:

```
key=value
```

For example, you could use the *key value* format to identify rules by author and date, using a category and sub-category as follows:

```
author SnortGuru_20050406
```

You can use multiple `metadata` keywords in a rule. You can also use commas to separate multiple *key value* statements in a single `metadata` keyword, as seen in the following example:

```
author SnortGuru_20050406, revised_by SnortUser1_20050707,
revised_by SnortUser2_20061003, revised_by
SnortUser1_20070123
```

You are not limited to using a *key value* or *key=value* format; however, you should be aware of limitations resulting from validation based on these formats.

Avoiding Restricted Characters

License: Protection

Note the following character restrictions:

- Do not use a semicolon (;) or colon (:) in a `metadata` keyword.
- Be aware when using commas that the system interprets a comma as a separator for multiple *key value* or *key=value* statements. For example:

  ```
  key value, key value, key value
  ```
- Be aware when using the equal to (=) character or space character that the system interprets these characters as separators between *key* and *value*. For example:

  ```
  key value
  key=value
  ```

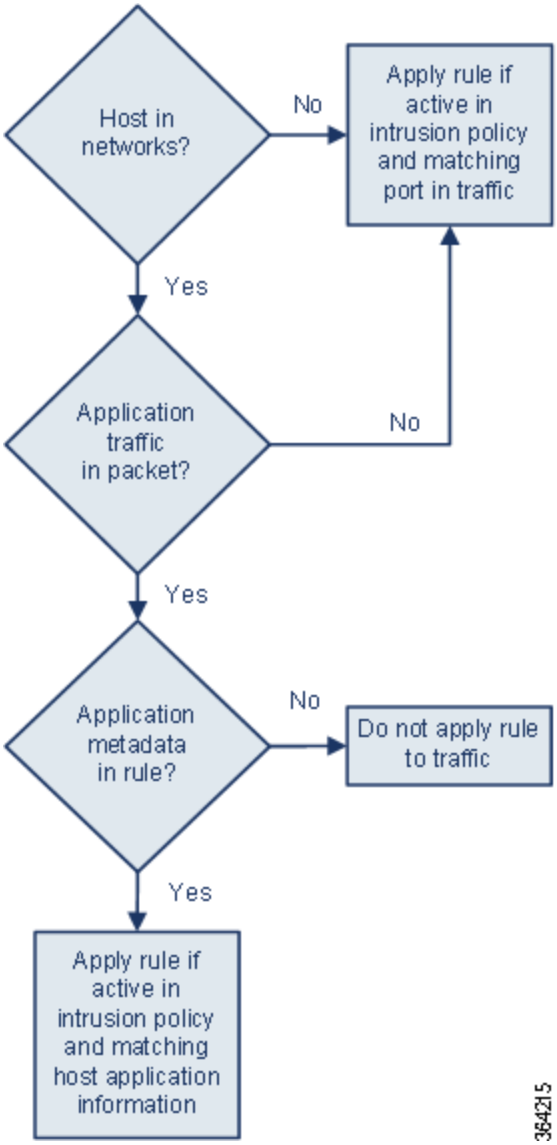
All other characters are permitted.

Adding service Metadata

License: Protection

The rules engine applies active rules with `service` metadata that match the application protocol information for the host in a packet to analyze and process traffic. If it does not match, the system does not apply the rule to the traffic. If a host does not have application protocol information, or if the rule does not have `service` metadata, the system checks the port in the traffic against the port in the rule to determine whether to apply the rule to the traffic.

The following diagram illustrates matching a rule to traffic based on application information:



To match a rule with an identified application protocol, you must define the `metadata` keyword and a *`key value`* statement, with `service` as the *`key`* and an application for the *`value`*. For example, the following *`key value`* statement in a `metadata` keyword associates the rule with HTTP traffic:

```
service http
```

The following table describes the most common application values.

Note Contact Support for assistance in defining applications not in the table.

Table 30-22 service Values

Value	Description
`dcerpc`	Distributed Computing Environment/Remote Procedure Calls System
`dns`	Domain Name System

Table 30-22 service Values (continued)

Value	Description
`finger`	Finger user information protocol
`ftp`	File Transfer Protocol
`ftp-data`	File Transfer Protocol (Data Channel)
`http`	Hypertext Transfer Protocol
`imap`	Internet Message Access Protocol
`isakmp`	Internet Security Association and Key Management Protocol
`netbios-dgm`	NETBIOS Datagram Service
`netbios-ns`	NETBIOS Name Service
`netbios-ssn`	NETBIOS Session Service
`nntp`	Network News Transfer Protocol
`oracle`	Oracle Net Services
`pop2`	Post Office Protocol, version 2
`pop3`	Post Office Protocol, version 3
`smtp`	Simple Mail Transfer Protocol
`ssh`	Secure Shell network protocol
`telnet`	Telnet network protocol
`tftp`	Trivial File Transfer Protocol
`x11`	X Window System

Avoiding Reserved Metadata

License: Protection

Avoid using the following words in a `metadata` keyword, either as a single argument or as the key in a *`key value`* statement; these are reserved for use by the VRT:

```
application
engine
impact_flag
os
policy
rule-type
rule-flushing
soid
```

Note Contact Support for assistance in adding restricted metadata to local rules that might not otherwise function as expected. See Importing Local Rule Files, page 46-14 for more information.

Inspecting IP Header Values

License: Protection

You can use keywords to identify possible attacks or security policy violations in the IP headers of packets. See the following sections for more information:

- Inspecting Fragments and Reserved Bits, page 30-46
- Inspecting the IP Header Identification Value, page 30-46
- Identifying Specified IP Options, page 30-47
- Identifying Specified IP Protocol Numbers, page 30-47
- Inspecting a Packet's Type of Service, page 30-47
- Inspecting a Packet's Time-To-Live Value, page 30-48

Inspecting Fragments and Reserved Bits

License: Protection

The `fragbits` keyword inspects the fragment and reserved bits in the IP header. You can check each packet for the Reserved Bit, the More Fragments bit, and the Don't Fragment bit in any combination.

Table 30-23 Fragbits Argument Values

Argument	Description
R	Reserved bit
M	More Fragments bit
D	Don't Fragment bit

To further refine a rule using the `fragbits` keyword, you can specify any operator described in the following table after the argument value in the rule.

Table 30-24 Fragbit Operators

Operator	Description
plus sign (`+`)	The packet must match against all specified bits.
asterisk (`*`)	The packet can match against any of the specified bits.
exclamation point (`!`)	The packet meets the criteria if none of the specified bits are set.

For example, to generate an event against packets that have the Reserved Bit set (and possibly any other bits), use `R+` as the `fragbits` value.

Inspecting the IP Header Identification Value

License: Protection

The `id` keyword tests the IP header fragment identification field against the value you specify in the keyword's argument. Some denial-of-service tools and scanners set this field to a specific number that is easy to detect. For example, in SID 630, which detects a Synscan portscan, the `id` value is set to `39426`, the static value used as the ID number in packets transmitted by the scanner.

Note `id` argument values must be numeric.

Identifying Specified IP Options

License: Protection

The `IPopts` keyword allows you to search packets for specified IP header options. The following table lists the available argument values.

Table 30-25 IPoption Arguments

Argument	Description
rr	record route
eol	end of list
nop	no operation
ts	time stamp
sec	IP security option
lsrr	loose source routing
ssrr	strict source routing
satid	stream identifier

Analysts most frequently watch for strict and loose source routing because these options may be an indication of a spoofed source IP address.

Identifying Specified IP Protocol Numbers

License: Protection

The `ip_proto` keyword allows you to identify packets with the IP protocol specified as the keyword's value. You can specify the IP protocols as a number, 0 through 255. You can find the complete list of protocol numbers at http://www.iana.org/assignments/protocol-numbers. You can combine these numbers with the following operators: `<`, `>`, or `!`. For example, to inspect traffic with any protocol that is not ICMP, use `!1` as a value to the `ip_proto` keyword. You can also use the `ip_proto` keyword multiple times in a single rule; note, however, that the rules engine interprets multiple instances of the keyword as having a Boolean AND relationship. For example, if you create a rule containing `ip_proto:!3; ip_proto:!6`, the rule ignores traffic using the GGP protocol AND the TCP protocol.

Inspecting a Packet's Type of Service

License: Protection

Some networks use the type of service (ToS) value to set precedence for packets traveling on that network. The `tos` keyword allows you to test the packet's IP header ToS value against the value you specify as the keyword's argument. Rules using the `tos` keyword will trigger on packets whose ToS is set to the specified value and that meet the rest of the criteria set forth in the rule.

Note Argument values for `tos` must be numeric.

The ToS field has been deprecated in the IP header protocol and replaced with the Differentiated Services Code Point (DSCP) field.

Inspecting a Packet's Time-To-Live Value

License: Protection

A packet's time-to-live (ttl) value indicates how many hops it can make before it is dropped. You can use the `ttl` keyword to test the packet's IP header ttl value against the value, or range of values, you specify as the keyword's argument. It may be helpful to set the `ttl` keyword parameter to a low value such as 0 or 1, as low time-to-live values are sometimes indicative of a traceroute or intrusion evasion attempt. (Note, though, that the appropriate value for this keyword depends on your device placement and network topology.) Use syntax as follows:

- Use an integer from 0 to 255 to set a specific value for the TTL value. You can also precede the value with an equal (=) sign (for example, you can specify `5` or `=5`).
- Use a hyphen (-) to specify a range of TTL values (for example, `0-2` specifies all values 0 through 2, `-5` specifies all values 0 through 5, and `5-` specifies all values 5 through 255).
- Use the greater than (>) sign to specify TTL values greater than a specific value (for example, `>3` specifies all values greater than 3).
- Use the greater than and equal to signs (>=) to specify TTL values greater than or equal to a specific value (for example, `>=3` specifies all values greater than or equal to 3).
- Use the less than (<) sign to specify TTL values less than a specific value (for example, `<3` specifies all values less than 3).
- Use the less than and equal to signs (<=) to specify TTL values less than or equal to a specific value (for example, `<=3` specifies all values less than or equal to 3).

Inspecting ICMP Header Values

License: Protection

The ASA FirePOWER module supports keywords that you can use to identify attacks and security policy violations in the headers of ICMP packets. Note, however, that predefined rules exist that detect most ICMP types and codes. Consider enabling an existing rule or creating a local rule based on an existing rule; you may be able to find a rule that meets your needs more quickly than if you build an ICMP rule from scratch.

See the following sections for more information about ICMP-specific keywords:

- Identifying Static ICMP ID and Sequence Values, page 30-48
- Inspecting the ICMP Message Type, page 30-49
- Inspecting the ICMP Message Code, page 30-49

Identifying Static ICMP ID and Sequence Values

License: Protection

The ICMP identification and sequence numbers help associate ICMP replies with ICMP requests. In normal traffic, these values are dynamically assigned to packets. Some covert channel and Distributed Denial of Server (DDoS) programs use static ICMP ID and sequence values. The following keywords allow you to identify ICMP packets with static values.

icmp_id

The `icmp_id` keyword inspects an ICMP echo request or reply packet's ICMP ID number. Use a numeric value that corresponds with the ICMP ID number as the argument for the `icmp_id` keyword.

icmp_seq

The `icmp_seq` keyword inspects an ICMP echo request or reply packet's ICMP sequence. Use a numeric value that corresponds with the ICMP sequence number as the argument for the `icmp_seq` keyword.

Inspecting the ICMP Message Type

License: Protection

Use the `itype` keyword to look for packets with specific ICMP message type values. You can specify either a valid ICMP type value (see http://www.iana.org/assignments/icmp-parameters or http://www.faqs.org/rfcs/rfc792.html for a full list of ICMP type numbers) or an invalid ICMP type value to test for different types of traffic. For example, attackers may set ICMP type values out of range to cause denial of service and flooding attacks.

You can specify a range for the `itype` argument value using less than (<) and greater than (>).

For example:

- <35
- >36
- 3<>55

Tip See http://www.iana.org/assignments/icmp-parameters or http://www.faqs.org/rfcs/rfc792.html for a full list of ICMP type numbers.

Inspecting the ICMP Message Code

License: Protection

ICMP messages sometimes include a code value that provides details when a destination is unreachable. (See the second section in http://www.iana.org/assignments/icmp-parameters for a full list of ICMP message codes correlated with the message types for which they can be used.)

You can use the `icode` keyword to identify packets with specific ICMP code values. You can choose to specify either a valid ICMP code value or an invalid ICMP code value to test for different types of traffic.

You can specify a range for the `icode` argument value using less than (<) and greater than (>).

For example:

- to find values less than 35, specify `<35`.
- to find values greater than 36, specify `>36`.
- to find values between 3 and 55, specify `3<>55`.

Tip You can use the `icode` and `itype` keywords together to identify traffic that matches both. For example, to identify ICMP traffic that contains an ICMP Destination Unreachable code type with an ICMP Port Unreachable code type, specify an `itype` keyword with a value of 3 (for Destination Unreachable) and an `icode` keyword with a value of 3 (for Port Unreachable).

Inspecting TCP Header Values and Stream Size

License: Protection

The ASA FirePOWER module supports keywords that are designed to identify attacks attempted using TCP headers of packets and TCP stream size. See the following sections for more information about TCP-specific keywords:

- Inspecting the TCP Acknowledgment Value, page 30-50
- Inspecting TCP Flag Combinations, page 30-50
- Applying Rules to a TCP or UDP Client or Server Flow, page 30-51
- Identifying Static TCP Sequence Numbers, page 30-52
- Identifying TCP Windows of a Given Size, page 30-53
- Identifying TCP Streams of a Given Size, page 30-53

Inspecting the TCP Acknowledgment Value

License: Protection

You can use the `ack` keyword to compare a value against a packet's TCP acknowledgment number. The rule triggers if a packet's TCP acknowledgment number matches the value specified for the `ack` keyword.

Argument values for `ack` must be numeric.

Inspecting TCP Flag Combinations

License: Protection

You can use the `flags` keyword to specify any combination of TCP flags that, when set in an inspected packet, cause the rule to trigger.

Note In situations where you would traditionally use `A+` as the value for `flags`, you should instead use the `flow` keyword with a value of `established`. Generally, you should use the `flow` keyword with a value of `stateless` when using flags to ensure that all combinations of flags are detected. See Applying Rules to a TCP or UDP Client or Server Flow, page 30-51 for more information about the `flow` keyword.

You can either check for or ignore the values described in the following table for the `flag` keyword.

Table 30-26 flag Arguments

Argument	TCP Flag
Ack	Acknowledges data.
Psh	Data should be sent in this packet.
Syn	A new connection.
Urg	Packet contains urgent data.
Fin	A closed connection.
Rst	An aborted connection.

Table 30-26 flag Arguments (continued)

Argument	TCP Flag
CWR	An ECN congestion window has been reduced. This was formerly the R1 argument, which is still supported for backward compatibility.
ECE	ECN echo. This was formerly the R2 argument, which is still supported for backward compatibility.

Tip For more information on Explicit Congestion Notification (ECN), see the information provided at: http://www.faqs.org/rfcs/rfc3168.html.

When using the `flags` keyword, you can use an operator to indicate how the system performs matches against multiple flags. The following table describes these operators.

Table 30-27 Operators Used with flags

Operator	Description	Example
all	The packet must contain all specified flags.	Select `Urg` and `all` to specify that a packet must contain the Urgent flag and may contain any other flags.
any	The packet can contain any of the specified flags.	Select `Ack`, `Psh`, and `any` to specify that either or both the `Ack` and `Psh` flags must be set to trigger the rule, and that other flags may also be set on a packet.
not	The packet must **not** contain the specified flag set.	Select `Urg` and `not` to specify that the Urgent flag is **not** set on packets that trigger this rule.

Applying Rules to a TCP or UDP Client or Server Flow

License: Protection

You can use the `flow` keyword to select packets for inspection by a rule based on session characteristics. The `flow` keyword allows you to specify the direction of the traffic flow to which a rule applies, applying rules to either the client flow or server flow. To specify how the `flow` keyword inspects your packets, you can set the direction of traffic you want analyzed, the state of packets inspected, and whether the packets are part of a rebuilt stream.

Stateful inspection of packets occurs when rules are processed. If you want a TCP rule to ignore stateless traffic (traffic without an established session context), you must add the `flow` keyword to the rule and select the **Established** argument for the keyword. If you want a UDP rule to ignore stateless traffic, you must add the `flow` keyword to the rule and select either the **Established** argument or a directional argument, or both. This causes the TCP or UDP rule to perform stateful inspection of a packet.

When you add a directional argument, the rules engine inspects only those packets that have an established state with a flow that matches the direction specified. For example, if you add the `flow` keyword with the `established` argument and the `From Client` argument to a rule that triggers when a TCP or UDP connection is detected, the rules engine only inspects packets that are sent from the client.

Tip For maximum performance, always include a `flow` keyword in a TCP rule or a UDP session rule.

To specify flow, select the `flow` keyword from the **Detection Options** list on the Create Rule page and click **Add Option**. Next, select the arguments from the list provided for each field.

The following table describes the stream-related arguments you can specify for the `flow` keyword:

Table 30-28 State-Related flow Arguments

Argument	Description
Established	Triggers on established connections.
Stateless	Triggers regardless of the state of the stream processor.

The following table describes the directional options you can specify for the `flow` keyword:

Table 30-29 flow Directional Arguments

Argument	Description
To Client	Triggers on server responses.
To Server	Triggers on client responses.
From Client	Triggers on client responses.
From Server	Triggers on server responses.

Notice that `From Server` and `To Client` perform the same function, as do `To Server` and `From Client`. These options exist to add context and readability to the rule. For example, if you create a rule designed to detect an attack from a server to a client, use `From Server`. But, if you create a rule designed to detect an attack from the client to the server, use `From Client`.

The following table describes the stream-related arguments you can specify for the `flow` keyword:

Table 30-30 Stream-Related flow Arguments

Argument	Description
Ignore Stream Traffic	Does not trigger on rebuilt stream packets.
Only Stream Traffic	Triggers only on rebuilt stream packets.

For example, you can use `To Server, Established, Only Stream Traffic` as the value for the `flow` keyword to detect traffic, traveling from a client to the server in an established session, that has been reassembled by the stream preprocessor.

Identifying Static TCP Sequence Numbers

License: Protection

The `seq` keyword allows you to specify a static sequence number value. Packets whose sequence number matches the specified argument trigger the rule containing the keyword. While this keyword is used rarely, it is helpful in identifying attacks and network scans that use generated packets with static sequence numbers.

Identifying TCP Windows of a Given Size

License: Protection

You can use the `window` keyword to specify the TCP window size you are interested in. A rule containing this keyword triggers whenever it encounters a packet with the specified TCP window size. While this keyword is used rarely, it is helpful in identifying attacks and network scans that use generated packets with static TCP window sizes.

Identifying TCP Streams of a Given Size

License: Protection

You can use the `stream_size` keyword in conjunction with the stream preprocessor to determine the size in bytes of a TCP stream, using the format:

`direction,operator,bytes`

where *`bytes`* is number of bytes. You must separate each option in the argument with a comma (,).

The following table describes the case-insensitive directional options you can specify for the `stream_size` keyword:

Table 30-31 stream_size Keyword Directional Arguments

Argument	Description
client	triggers on a stream from the client matching the specified stream size.
server	triggers on a stream from the server matching the specified stream size.
both	triggers on traffic from the client and traffic from the server both matching the specified stream size. For example, the argument `both, >, 200` would trigger when traffic from the client is greater than 200 bytes AND traffic from the server is greater than 200 bytes.
either	triggers on traffic from either the client or the server matching the specified stream size, whichever occurs first. For example, the argument `either, >, 200` would trigger when traffic from the client is greater than 200 bytes OR traffic from the server is greater than 200 bytes.

The following table describes the operators you can use with the `stream_size` keyword:

Table 30-32 stream_size Keyword Argument Operators

Operator	Description
=	equal to
!=	not equal to
>	greater than
<	less than
>=	greater than or equal to
<=	less than or equal to

For example, you could use `client, >=, 5001216` as the argument for the `stream_size` keyword to detect a TCP stream traveling from a client to a server and greater than or equal to 5001216 bytes.

Enabling and Disabling TCP Stream Reassembly

License: Protection

You can use the `stream_reassemble` keyword to enable or disable TCP stream reassembly for a single connection when inspected traffic on the connection matches the conditions of the rule. Optionally, you can use this keyword multiple times in a rule.

Use the following syntax to enable or disable stream reassembly:

```
enable|disable, server|client|both, option, option
```

The following table describes the optional arguments you can use with the `stream_reassemble` keyword.

Table 30-33 stream_reassemble Optional Arguments

Argument	Description
noalert	Generate no events regardless of any other detection options specified in the rule.
fastpath	Ignore the rest of the connection traffic when there is a match.

For example, the following rule disables TCP client-side stream reassembly without generating an event on the connection where a 200 OK status code is detected in an HTTP response:

```
alert tcp any 80 -> any any (flow:to_client, established; content: "200 OK";
stream_reassemble:disable, client, noalert
```

To use stream_reassemble:

Step 1 On the Create Rule page, select `stream_reassemble` in the drop-down list and click **Add Option**.

The `stream_reassemble` section appears.

Extracting SSL Information from a Session

License: Protection

You can use SSL rule keywords to invoke the Secure Sockets Layer (SSL) preprocessor and extract information about SSL version and session state from packets in an encrypted session.

When a client and server communicate to establish an encrypted session using SSL or Transport Layer Security (TLS), they exchange handshake messages. Although the data transmitted in the session is encrypted, the handshake messages are not.

The SSL preprocessor extracts state and version information from specific handshake fields. Two fields within the handshake indicate the version of SSL or TLS used to encrypt the session and the stage of the handshake.

For more information, see the following sections:

- ssl_state, page 30-55
- ssl_version, page 30-55

ssl_state

License: Protection

The `ssl_state` keyword can be used to match against state information for an encrypted session. To check for two or more SSL versions used simultaneously, use multiple `ssl_version` keywords in a rule.

When a rule uses the `ssl_state` keyword, the rules engine invokes the SSL preprocessor to check traffic for SSL state information.

For example, to detect an attacker's attempt to cause a buffer overflow on a server by sending a `ClientHello` message with an overly long challenge length and too much data, you could use the `ssl_state` keyword with `client_hello` as an argument then check for abnormally large packets.

Use a comma-separated list to specify multiple arguments for the SSL state. When you list multiple arguments, the system evaluates them using the OR operator. For example, if you specify `client_hello` and `server_hello` as arguments, the system evaluates the rule against traffic that has a `client_hello` OR a `server_hello`.

You can also negate any argument; for example:

```
!client_hello, !unknown
```

To ensure the connection has reached each of a set of states, multiple rules using the ssl_state rule option should be used. The `ssl_state` keyword takes the following identifiers as arguments:

Table 30-34 `ssl_state` ***Arguments***

Argument	Purpose
`client_hello`	Matches against a handshake message with `ClientHello` as the message type, where the client requests an encrypted session.
`server_hello`	Matches against a handshake message with `ServerHello` as the message type, where the server responds to the client's request for an encrypted session.
`client_keyx`	Matches against a handshake message with `ClientKeyExchange` as the message type, where the client transmits a key to the server to confirm receipt of a key from the server.
`server_keyx`	Matches against a handshake message with `ServerKeyExchange` as the message type, where the client transmits a key to the server to confirm receipt of a key from the server.
`unknown`	Matches against any handshake message type.

ssl_version

License: Protection

The `ssl_version` keyword can be used to match against version information for an encrypted session. When a rule uses the `ssl_version` keyword, the rules engine invokes the SSL preprocessor to check traffic for SSL version information.

For example, if you know there is a buffer overflow vulnerability in SSL version 2, you could use the `ssl_version` keyword with the `sslv2` argument to identify traffic using that version of SSL.

Use a comma-separated list to specify multiple arguments for the SSL version. When you list multiple arguments, the system evaluates them using the OR operator. For example, if you wanted to identify any encrypted traffic that was not using SSLv2, you could add `ssl_version:ssl_v3,tls1.0,tls1.1,tls1.2` to a rule. The rule would evaluate any traffic using SSL Version 3, TLS Version 1.0, TLS Version 1.1, or TLS Version 1.2.

The `ssl_version` keyword takes the following SSL/TLS version identifiers as arguments:

Table 30-35 `ssl_version` Arguments

Argument	Purpose
`sslv2`	Matches against traffic encoded using Secure Sockets Layer (SSL) Version 2.
`sslv3`	Matches against traffic encoded using Secure Sockets Layer (SSL) Version 3.
`tls1.0`	Matches against traffic encoded using Transport Layer Security (TLS) Version 1.0.
`tls1.1`	Matches against traffic encoded using Transport Layer Security (TLS) Version 1.1.
`tls1.2`	Matches against traffic encoded using Transport Layer Security (TLS) Version 1.2.

Inspecting Application Layer Protocol Values

License: Protection

Although preprocessors perform most of the normalization and inspection of application layer protocol values, you can continue to inspect application layer values using the keywords described in the following sections:

- RPC, page 30-56
- ASN.1, page 30-57
- urilen, page 30-58
- DCE/RPC Keywords, page 30-59
- SIP Keywords, page 30-62
- GTP Keywords, page 30-64
- Modbus Keywords, page 30-74
- DNP3 Keywords, page 30-76

RPC

License: Protection

The `rpc` keyword identifies Open Network Computing Remote Procedure Call (ONC RPC) services in TCP or UDP packets. This allows you to detect attempts to identify the RPC programs on a host. Intruders can use an RPC portmapper to determine if any of the RPC services running on your network can be exploited. They can also attempt to access other ports running RPC without using portmapper. The following table lists the arguments that the `rpc` keyword accepts.

Table 30-36 rpc Keyword Arguments

Argument	Description
application	The RPC application number
procedure	The RPC procedure invoked
version	The RPC version

To specify the arguments for the `rpc` keyword, use the following syntax:

application,*procedure*,*version*

where *application* is the RPC application number, *procedure* is the RPC procedure number, and *version* is the RPC version number. You must specify all arguments for the `rpc` keyword — if you are not able to specify one of the arguments, replace it with an asterisk (`*`).

For example, to search for RPC portmapper (which is the RPC application indicated by the number 100000), with any procedure or version, use `100000,*,*` as the arguments.

ASN.1

License: Protection

The `asn1` keyword allows you to decode a packet or a portion of a packet, looking for various malicious encodings.

The following table describes the arguments for the `asn1` keyword.

Table 30-37 asn.1 Keyword Arguments

Argument	Description
Bitstring Overflow	Detects invalid, remotely exploitable bitstring encodings.
Double Overflow	Detects a double ASCII encoding that is larger than a standard buffer. This is known to be an exploitable function in Microsoft Windows, but it is unknown at this time which services may be exploitable.
Oversize Length	Detects ASN.1 type lengths greater than the supplied argument. For example, if you set the Oversize Length to 500, any ASN.1 type greater than 500 triggers the rule.
Absolute Offset	Sets an absolute offset from the beginning of the packet payload. (Remember that the offset counter starts at byte 0.) For example, if you want to decode SNMP packets, set Absolute Offset to 0 and do not set a Relative Offset. Absolute Offset may be positive or negative.
Relative Offset	This is the relative offset from the last successful content match, `pcre`, or `byte_jump`. To decode an ASN.1 sequence right after the content "foo", set Relative Offset to 0, and do not set an Absolute Offset. Relative Offset may be positive or negative. (Remember that the offset counter starts at 0.)

For example, there is a known vulnerability in the Microsoft ASN.1 Library that creates a buffer overflow, allowing an attacker to exploit the condition with a specially crafted authentication packet. When the system decodes the asn.1 data, exploit code in the packet could execute on the host with system-level privileges or could cause a DoS condition. The following rule uses the `asn1` keyword to detect attempts to exploit this vulnerability:

```
alert tcp $EXTERNAL_NET any -> $HOME_NET 445
(flow:to_server, established; content:"|FF|SMB|73|"; nocase;
offset:4; depth:5;
asn1:bitstring_overflow,double_overflow,oversize_length
100,relative_offset 54;)
```

The above rule generates an event against TCP traffic traveling from any IP address defined in the $EXTERNAL_NET variable, from any port, to any IP address defined in the $HOME_NET variable using port 445. In addition, it only executes the rule on established TCP connections to servers. The rule then tests for specific content in specific locations. Finally, the rule uses the `asn1` keyword to detect

bitstring encodings and double ASCII encodings and to identify asn.1 type lengths over 100 bytes in length starting 55 bytes from the end of the last successful content match. (Remember that the `offset` counter starts at byte 0.)

urilen

License: Protection

You can use the `urilen` keyword in conjunction with the HTTP Inspect preprocessor to inspect HTTP traffic for URIs of a specific length, less than a maximum length, greater than a minimum length, or within a specified range.

After the HTTP Inspect preprocessor normalizes and inspects the packet, the rules engine evaluates the packet against the rule and determines whether the URI matches the length condition specified by the `urilen` keyword. You can use this keyword to detect exploits that attempt to take advantage of URI length vulnerabilities, for example, by creating a buffer overflow that allows the attacker to cause a DoS condition or execute code on the host with system-level privileges.

Note the following when using the `urilen` keyword in a rule:

- In practice, you always use the `urilen` keyword in combination with the `flow:established` keyword and one or more other keywords.
- The rule protocol is always TCP. See Specifying Protocols, page 30-4 for more information.
- Target ports are always HTTP ports. See Defining Ports in Intrusion Rules, page 30-8 and Optimizing Predefined Default Variables, page 2-13 for more information.

You specify the URI length using a decimal number of bytes, less than (<) and greater than (>).

For example:

- specify `5` to detect a URI 5 bytes long.
- specify `< 5` (separated by one space character) to detect a URI less than 5 bytes long.
- specify `> 5` (separated by one space character) to detect a URI greater than 5 bytes long.
- specify `3 <> 5` (with one space character before and after `<>`) to detect a URI between 3 and 5 bytes long inclusive.

For example, there is a known vulnerability in Novell's server monitoring and diagnostics utility iMonitor version 2.4, which comes with eDirectory version 8.8. A packet containing an excessively long URI creates a buffer overflow, allowing an attacker to exploit the condition with a specially crafted packet that could execute on the host with system-level privileges or could cause a DoS condition. The following rule uses the `urilen` keyword to detect attempts to exploit this vulnerability:

```
alert tcp $EXTERNAL_NET any -> $HOME_NET $HTTP_PORTS
(msg:"EXPLOIT eDirectory 8.8 Long URI iMonitor buffer
overflow attempt";flow:to_server,established;
urilen:> 8192; uricontent:"/nds/"; nocase;
classtype:attempted-admin; sid:x; rev:1;)
```

The above rule generates an event against TCP traffic traveling from any IP address defined in the $EXTERNAL_NET variable, from any port, to any IP address defined in the $HOME_NET variable using the ports defined in the $HTTP_PORTS variable. In addition, packets are evaluated against the rule only on established TCP connections to servers. The rule uses the `urilen` keyword to detect any URI over 8192 bytes in length. Finally, the rule searches the URI for the specific case-insensitive content `/nds/`.

DCE/RPC Keywords

License: Protection

The three DCE/RPC keywords described in the following table allow you to monitor DCE/RPC session traffic for exploits. When the system processes rules with these keywords, it invokes the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic, page 22-2 for more information.

Table 30-38 DCE/RPC Keywords

Use...	In this way...	To detect...
`dce_iface`	alone	packets identifying a specific DCE/RPC service
`dce_opnum`	preceded by `dce_iface`	packets identifying specific DCE/RPC service operations
`dce_stub_data`	preceded by `dce_iface` + `dce_opnum`	stub data defining a specific operation request or response

Note in the table that you should always precede `dce_opnum` with `dce_iface`, and you should always precede `dce_stub_data` with `dce_iface` + `dce_opnum`.

You can also use these DCE/RPC keywords in combination with other rule keywords. Note that for DCE/RPC rules, you use the `byte_jump`, `byte_test`, `and byte_extract` keywords with their **DCE/RPC** arguments selected. For more information, see Using Byte_Jump and Byte_Test, page 30-30 and Reading Packet Data into Keyword Arguments, page 30-82.

Cisco recommends that you include at least one `content` keyword in rules that include DCE/RPC keywords to ensure that the rules engine uses the fast pattern matcher, which increases processing speed and improves performance. Note that the rules engine uses the fast pattern matcher when a rule includes at least one `content` keyword, regardless of whether you enable the `content` keyword **Use Fast Pattern Matcher** argument. See Searching for Content Matches, page 30-15 and Use Fast Pattern Matcher, page 30-26 for more information.

You can use the DCE/RPC version and adjoining header information as the matching content in the following cases:

- the rule does not include another `content` keyword
- the rule contains another `content` keyword, but the DCE/RPC version and adjoining information represent a more unique pattern than the other content

  For example, the DCE/RPC version and adjoining information are more likely to be unique than a single byte of content.

You should end qualifying rules with one of the following version and adjoining information content matches:

- For connection-oriented DCE/RPC rules, use the content `|05 00 00|` (for major version 05, minor version 00, and the request PDU (protocol data unit) type 00).
- For connectionless DCE/RPC rules, use the content `|04 00|` (for version 04, and the request PDU type 00).

In either case, position the `content` keyword for version and adjoining information as the last keyword in the rule to invoke the fast pattern matcher without repeating processing already completed by the DCE/RPC preprocessor. Note that placing the `content` keyword at the end of the rule applies to version content used as a device to invoke the fast pattern matcher, and not necessarily to other content matches in the rule.

See the following sections for more information:

- dce_iface, page 30-60
- dce_opnum, page 30-61
- dce_stub_data, page 30-61

dce_iface

License: Protection

You can use the `dce_iface` keyword to identify a specific DCE/RPC service.

Optionally, you can also use `dce_iface` in combination with the `dce_opnum` and `dce_stub_data` keywords to further limit the DCE/RPC traffic to inspect. See dce_opnum, page 30-61 and dce_stub_data, page 30-61 for more information.

A fixed, sixteen-byte Universally Unique Identifier (UUID) identifies the application interface assigned to each DCE/RPC service. For example, the UUID 4b324fc8-670-01d3-1278-5a47bf6ee188 identifies the DCE/RPC lanmanserver service, also known as the srvsvc service, which provides numerous management functions for sharing peer-to-peer printers, files, and SMB named pipes. The DCE/RPC preprocessor uses the UUID and associated header values to track DCE/RPC sessions.

The interface UUID is comprised of five hexadecimal strings separated by hyphens:

```
<4hexbytes>-<2hexbytes>-<2hexbytes>-<2hexbytes>-<6hexbytes>
```

You specify the interface by entering the entire UUID including hyphens, as seen in the following UUID for the netlogon interface:

```
12345678-1234-abcd-ef00-01234567cffb
```

Note that you must specify the first three strings in the UUID in big endian byte order. Although published interface listings and protocol analyzers typically display UUIDs in the correct byte order, you might encounter a need to rearrange the UUID byte order before entering it. Consider the following messenger service UUID shown as it might sometimes be displayed in raw ASCII text with the first three strings in little endian byte order:

```
f8 91 7b 5a 00 ff d0 11 a9 b2 00 c0 4f b6 e6 fc
```

You would specify the same UUID for the `dce_iface` keyword by inserting hyphens and putting the first three strings in big endian byte order as follows:

```
5a7b91f8-ff00-11d0-a9b2-00c04fb6e6fc
```

Although a DCE/RPC session can include requests to multiple interfaces, you should include only one `dce_iface` keyword in a rule. Create additional rules to detect additional interfaces.

DCE/RPC application interfaces also have interface version numbers. You can optionally specify an interface version with an operator indicating that the version equals, does not equal, is less than, or greater than the specified value.

Both connection-oriented and connectionless DCE/RPC can be fragmented in addition to any TCP segmentation or IP fragmentation. Typically, it is not useful to associate any DCE/RPC fragment other than the first with the specified interface, and doing so may result in a large number of false positives. However, for flexibility you can optionally evaluate all fragments against the specified interface.

The following table summarizes the `dce_iface` keyword arguments.

Table 30-39 `dce_iface` *Arguments*

Argument	Description
Interface UUID	The UUID, including hyphens, that identifies the application interface of the specific service that you want to detect in DCE/RPC traffic. Any request associated with the specified interface would match the interface UUID.
Version	Optionally, the application interface version number 0 to 65535 and an operator indicating whether to detect a version greater than (>), less than (<), equal to (=), or not equal to (!) the specified value.
All Fragments	Optionally, enable to match against the interface in all associated DCE/RPC fragments and, if specified, on the interface version. This argument is disabled by default, indicating that the keyword matches only if the first fragment or the entire unfragmented packet is associated with the specified interface. Note that enabling this argument may result in false positives.

dce_opnum

License: Protection

You can use the `dce_opnum` keyword in conjunction with the DCE/RPC preprocessor to detect packets that identify one or more specific operations that a DCE/RPC service provides.

Client function calls request specific service functions, which are referred to in DCE/RPC specifications as *operations*. An operation number (opnum) identifies a specific operation in the DCE/RPC header. It is likely that an exploit would target a specific operation.

For example, the UUID 12345678-1234-abcd-ef00-01234567cffb identifies the interface for the netlogon service, which provides several dozen different operations. One of these is operation 6, the NetrServerPasswordSet operation.

You should precede a `dce_opnum` keyword with a `dce_iface` keyword to identify the service for the operation. See dce_iface, page 30-60 for more information.

You can specify a single decimal value 0 to 65535 for a specific operation, a range of operations separated by a hyphen, or a comma-separated list of operations and ranges in any order.

Any of the following examples would specify valid netlogon operation numbers:

```
15
15-18
15, 18-20
15, 20-22, 17
15, 18-20, 22, 24-26
```

dce_stub_data

License: Protection

You can use the `dce_stub_data` keyword in conjunction with the DCE/RPC preprocessor to specify that the rules engine should start inspection at the beginning of the stub data, regardless of any other rule options. Packet payload rule options that follow the `dce_stub_data` keyword are applied relative to the stub data buffer.

DCE/RPC stub data provides the interface between a client procedure call and the DCE/RPC run-time system, the mechanism that provides the routines and services central to DCE/RPC. DCE/RPC exploits are identified in the stub data portion of the DCE/RPC packet. Because stub data is associated with a specific operation or function call, you should always precede `dce_stub_data` with `dce_iface` and `dce_opnum` to identify the related service and operation.

The `dce_stub_data` keyword has no arguments. See dce_iface, page 30-60 and dce_opnum, page 30-61 for more information.

SIP Keywords

License: Protection

Four SIP keywords allow you to monitor SIP session traffic for exploits.

Note that the SIP protocol is vulnerable to denial of service (DoS) attacks. Rules addressing these attacks can benefit from rate-based attack prevention. See Adding Dynamic Rule States, page 27-28 and Preventing Rate-Based Attacks, page 28-9 for more information.

See the following sections for more information:

- sip_header, page 30-62
- sip_body, page 30-62
- sip_method, page 30-62
- sip_stat_code, page 30-63

sip_header

License: Protection

You can use the `sip_header` keyword to start inspection at the beginning of the extracted SIP request or response header and restrict inspection to header fields.

The `sip_header` keyword has no arguments. See sip_method, page 30-62 and sip_stat_code, page 30-63 for more information.

The following example rule fragment points to the SIP header and matches the CSeq header field:

```
alert udp any any -> any 5060 ( sip_header; content:"CSeq"; )
```

sip_body

License: Protection

You can use the `sip_body` keyword to start inspection at the beginning of the extracted SIP request or response message body and restrict inspection to the message body.

The `sip_body` keyword has no arguments.

The following example rule fragment points to the SIP message body and matches a specific IP address in the c (connection information) field in extracted SDP data:

```
alert udp any any -> any 5060 ( sip_body; content:"c=IN 192.168.12.14"; )
```

Note that rules are not limited to searching for SDP content. The SIP preprocessor extracts the entire message body and makes it available to the rules engine.

sip_method

License: Protection

A *method* field in each SIP request identifies the purpose of the request. You can use the `sip_method` keyword to test SIP requests for specific methods. Separate multiple methods with commas.

You can specify any of the following currently defined SIP methods:

```
ack, benotify, bye, cancel, do, info, invite, join, message, notify, options, prack,
publish, quath, refer, register, service, sprack, subscribe, unsubscribe, update
```

Methods are case-insensitive. You can separate multiple methods with commas.

Because new SIP methods might be defined in the future, you can also specify a custom method, that is, a method that is not a currently defined SIP method. Accepted field values are defined in RFC 2616, which allows all characters except control characters and separators such as `=`, `(`, and `}`. See RFC 2616 for the complete list of excluded separators. When the system encounters a specified custom method in traffic, it will inspect the packet header but not the message.

The system supports up to 32 methods, including the 21 currently defined methods and an additional 11 methods. The system ignores any undefined methods that you might configure. Note that the 32 total methods includes methods specified using the **Methods to Check** SIP preprocessor option. See Selecting SIP Preprocessor Options, page 22-49 for more information.

You can specify only one method when you use negation. For example:

```
!invite
```

Note, however, that multiple `sip_method` keywords in a rule are linked with an **AND** operation. For example, to test for all extracted methods except `invite` and `cancel`, you would use two negated `sip_method` keywords:

```
sip_method: !invite
sip_method: !cancel
```

Cisco recommends that you include at least one `content` keyword in rules that include the `sip_method` keyword to ensure that the rules engine uses the fast pattern matcher, which increases processing speed and improves performance. Note that the rules engine uses the fast pattern matcher when a rule includes at least one `content` keyword, regardless of whether you enable the `content` keyword **Use Fast Pattern Matcher** argument. See Searching for Content Matches, page 30-15 and Use Fast Pattern Matcher, page 30-26 for more information.

sip_stat_code

License: Protection

A three-digit status code in each SIP response indicates the outcome of the requested action. You can use the `sip_stat_code` keyword to test SIP responses for specific status codes.

You can specify a one-digit response-type number 1-9, a specific three-digit number 100-999, or a comma-separated list of any combination of either. A list matches if any single number in the list matches the code in the SIP response.

The following table describes the SIP status code values you can specify.

Table 30-40 `sip_stat_code` ***Values***

To detect...	Specify...	For example...	Detects...
a specific status code	the three-digit status code	189	189

Table 30-40 `sip_stat_code` *Values (continued)*

To detect...	Specify...	For example...	Detects...
any three-digit code that begins with a specified single digit	the single digit	1	1xx; that is, 100, 101, 102, and so on
a list of values	any comma-separated combination of specific codes and single digits	222, 3	222 plus 300, 301, 302, and so on

Note also that the rules engine does not use the fast pattern matcher to search for the value specify using the `sip_stat_code` keyword, regardless of whether your rule includes a `content` keyword.

GTP Keywords

License: Protection

Three GSRP Tunneling Protocol (GTP) keywords allow you to inspect the GTP command channel for GTP version, message type, and information elements. You cannot use GTP keywords in combination with other intrusion rule keywords such as `content` or `byte_jump`. You must use the `gtp_version` keyword in each rule that uses the `gtp_info` or `gtp_type` keyword.

See the following sections for more information:

- gtp_version, page 30-64
- gtp_type, page 30-64
- gtp_info, page 30-69

gtp_version

You can use the `gtp_version` keyword to inspect GTP control messages for GTP version 0, 1, or 2.

Because different GTP versions define different message types and information elements, you must use this keyword when you use the `gtp_type` or `gtp_info` keyword. You can specify the value 0, 1, or 2.

To specify the GTP version:

Step 1 On the Create Rule page, select **gtp_version** in the drop-down list and click **Add Option.**

The `gtp_version` keyword appears.

Step 2 Specify `0`, `1`, or `2` to identify the GTP version.

gtp_type

Each GTP message is identified by a message type, which is comprised of both a numeric value and a string. You can use the `gtp_type` keyword in combination with the `gtp_version` keyword to inspect traffic for specific GTP message types.

You can specify a defined decimal value for a message type, a defined string, or a comma-separated list of either or both in any combination, as seen in the following example:

```
10, 11, echo_request
```

The system uses an OR operation to match each value or string that you list. The order in which you list values and strings does not matter. Any single value or string in the list matches the keyword. You receive an error if you attempt to save a rule that includes an unrecognized string or an out-of-range value.

Note in the table that different GTP versions sometimes use different values for the same message type. For example, the `sgsn_context_request` message type has a value of 50 in GTPv0 and GTPv1, but a value of 130 in GTPv2.

The `gtp_type` keyword matches different values depending on the version number in the packet. In the example above, the keyword matches the message type value 50 in a GTPv0 or GTPv1 packet and the value 130 in a GTPv2 packet. The keyword does not match a packet when the message type value in the packet is not a known value for the version specified in the packet.

If you specify an integer for the message type, the keyword matches if the message type in the keyword matches the value in the GTP packet, regardless of the version specified in the packet.

The following table lists the defined values and strings recognized by the system for each GTP message type.

Table 30-41 GTP Message Types

Value	Version 0	Version 1	Version 2
1	echo_request	echo_request	echo_request
2	echo_response	echo_response	echo_response
3	version_not_supported	version_not_supported	version_not_supported
4	node_alive_request	node_alive_request	N/A
5	node_alive_response	node_alive_response	N/A
6	redirection_request	redirection_request	N/A
7	redirection_response	redirection_response	N/A
16	create_pdp_context_request	create_pdp_context_request	N/A
17	create_pdp_context_response	create_pdp_context_response	N/A
18	update_pdp_context_request	update_pdp_context_request	N/A
19	update_pdp_context_response	update_pdp_context_response	N/A
20	delete_pdp_context_request	delete_pdp_context_request	N/A
21	delete_pdp_context_response	delete_pdp_context_response	N/A
22	create_aa_pdp_context_request	init_pdp_context_activation_request	N/A
23	create_aa_pdp_context_response	init_pdp_context_activation_response	N/A
24	delete_aa_pdp_context_request	N/A	N/A
25	delete_aa_pdp_context_response	N/A	N/A
26	error_indication	error_indication	N/A
27	pdu_notification_request	pdu_notification_request	N/A
28	pdu_notification_response	pdu_notification_response	N/A
29	pdu_notification_reject_request	pdu_notification_reject_request	N/A
30	pdu_notification_reject_response	pdu_notification_reject_response	N/A

Table 30-41 GTP Message Types (continued)

Value	Version 0	Version 1	Version 2
31	N/A	supported_ext_header_notification	N/A
32	send_routing_info_request	send_routing_info_request	create_session_request
33	send_routing_info_response	send_routing_info_response	create_session_response
34	failure_report_request	failure_report_request	modify_bearer_request
35	failure_report_response	failure_report_response	modify_bearer_response
36	note_ms_present_request	note_ms_present_request	delete_session_request
37	note_ms_present_response	note_ms_present_response	delete_session_response
38	N/A	N/A	change_notification_request
39	N/A	N/A	change_notification_response
48	identification_request	identification_request	N/A
49	identification_response	identification_response	N/A
50	sgsn_context_request	sgsn_context_request	N/A
51	sgsn_context_response	sgsn_context_response	N/A
52	sgsn_context_ack	sgsn_context_ack	N/A
53	N/A	forward_relocation_request	N/A
54	N/A	forward_relocation_response	N/A
55	N/A	forward_relocation_complete	N/A
56	N/A	relocation_cancel_request	N/A
57	N/A	relocation_cancel_response	N/A
58	N/A	forward_srns_contex	N/A
59	N/A	forward_relocation_complete_ack	N/A
60	N/A	forward_srns_contex_ack	N/A
64	N/A	N/A	modify_bearer_command
65	N/A	N/A	modify_bearer_failure_indication
66	N/A	N/A	delete_bearer_command
67	N/A	N/A	delete_bearer_failure_indication
68	N/A	N/A	bearer_resource_command
69	N/A	N/A	bearer_resource_failure_indication
70	N/A	ran_info_relay	downlink_failure_indication
71	N/A	N/A	trace_session_activation
72	N/A	N/A	trace_session_deactivation
73	N/A	N/A	stop_paging_indication
95	N/A	N/A	create_bearer_request
96	N/A	mbms_notification_request	create_bearer_response
97	N/A	mbms_notification_response	update_bearer_request

Table 30-41 GTP Message Types (continued)

Value	Version 0	Version 1	Version 2
98	N/A	mbms_notification_reject_request	update_bearer_response
99	N/A	mbms_notification_reject_response	delete_bearer_request
100	N/A	create_mbms_context_request	delete_bearer_response
101	N/A	create_mbms_context_response	delete_pdn_request
102	N/A	update_mbms_context_request	delete_pdn_response
103	N/A	update_mbms_context_response	N/A
104	N/A	delete_mbms_context_request	N/A
105	N/A	delete_mbms_context_response	N/A
112	N/A	mbms_register_request	N/A
113	N/A	mbms_register_response	N/A
114	N/A	mbms_deregister_request	N/A
115	N/A	mbms_deregister_response	N/A
116	N/A	mbms_session_start_request	N/A
117	N/A	mbms_session_start_response	N/A
118	N/A	mbms_session_stop_request	N/A
119	N/A	mbms_session_stop_response	N/A
120	N/A	mbms_session_update_request	N/A
121	N/A	mbms_session_update_response	N/A
128	N/A	ms_info_change_request	identification_request
129	N/A	ms_info_change_response	identification_response
130	N/A	N/A	sgsn_context_request
131	N/A	N/A	sgsn_context_response
132	N/A	N/A	sgsn_context_ack
133	N/A	N/A	forward_relocation_request
134	N/A	N/A	forward_relocation_response
135	N/A	N/A	forward_relocation_complete
136	N/A	N/A	forward_relocation_complete_ack
137	N/A	N/A	forward_access
138	N/A	N/A	forward_access_ack
139	N/A	N/A	relocation_cancel_request
140	N/A	N/A	relocation_cancel_response
141	N/A	N/A	configuration_transfer_tunnel
149	N/A	N/A	detach
150	N/A	N/A	detach_ack
151	N/A	N/A	cs_paging

Table 30-41 GTP Message Types (continued)

Value	Version 0	Version 1	Version 2
152	N/A	N/A	ran_info_relay
153	N/A	N/A	alert_mme
154	N/A	N/A	alert_mme_ack
155	N/A	N/A	ue_activity
156	N/A	N/A	ue_activity_ack
160	N/A	N/A	create_forward_tunnel_request
161	N/A	N/A	create_forward_tunnel_response
162	N/A	N/A	suspend
163	N/A	N/A	suspend_ack
164	N/A	N/A	resume
165	N/A	N/A	resume_ack
166	N/A	N/A	create_indirect_forward_tunnel_request
167	N/A	N/A	create_indirect_forward_tunnel_response
168	N/A	N/A	delete_indirect_forward_tunnel_request
169	N/A	N/A	delete_indirect_forward_tunnel_response
170	N/A	N/A	release_access_bearer_request
171	N/A	N/A	release_access_bearer_response
176	N/A	N/A	downlink_data
177	N/A	N/A	downlink_data_ack
179	N/A	N/A	pgw_restart
180	N/A	N/A	pgw_restart_ack
200	N/A	N/A	update_pdn_request
201	N/A	N/A	update_pdn_response
211	N/A	N/A	modify_access_bearer_request
212	N/A	N/A	modify_access_bearer_response
231	N/A	N/A	mbms_session_start_request
232	N/A	N/A	mbms_session_start_response
233	N/A	N/A	mbms_session_update_request
234	N/A	N/A	mbms_session_update_response
235	N/A	N/A	mbms_session_stop_request
236	N/A	N/A	mbms_session_stop_response
240	data_record_transfer_request	data_record_transfer_request	N/A
241	data_record_transfer_response	data_record_transfer_response	N/A
254	N/A	end_marker	N/A
255	pdu	pdu	N/A

To specify GTP message types:

Step 1 On the Create Rule page, select **gtp_type** in the drop-down list and click **Add Option.**

The `gtp_type` keyword appears.

Step 2 Specify a defined decimal value 0 to 255 for the message type, a defined string, or a comma-separated list of either or both in any combination. See the GTP Message Types table for values and strings recognized by the system.

gtp_info

A GTP message can include multiple information elements, each of which is identified by both a defined numeric value and a defined string. You can use the `gtp_info` keyword in combination with the `gtp_version` keyword to start inspection at the beginning of a specified information element and restrict inspection to the specified information element.

You can specify either the defined decimal value or the defined string for an information element. You can specify a single value or string, and you can use multiple `gtp_info` keywords in a rule to inspect multiple information elements.

When a message includes multiple information elements of the same type, all are inspected for a match. When information elements occur in an invalid order, only the last instance is inspected.

Note that different GTP versions sometimes use different values for the same information element. For example, the `cause` information element has a value of 1 in GTPv0 and GTPv1, but a value of 2 in GTPv2.

The `gtp_info` keyword matches different values depending on the version number in the packet. In the example above, the keyword matches the information element value 1 in a GTPv0 or GTPv1 packet and the value 2 in a GTPv2 packet. The keyword does not match a packet when the information element value in the packet is not a known value for the version specified in the packet.

If you specify an integer for the information element, the keyword matches if the message type in the keyword matches the value in the GTP packet, regardless of the version specified in the packet.

The following table lists the values and strings recognized by the system for each GTP information element.

Table 30-42 GTP Information Elements

Value	Version 0	Version 1	Version 2
1	cause	cause	imsi
2	imsi	imsi	cause
3	rai	rai	recovery
4	tlli	tlli	N/A
5	p_tmsi	p_tmsi	N/A
6	qos	N/A	N/A
8	recording_required	recording_required	N/A
9	authentication	authentication	N/A
11	map_cause	map_cause	N/A
12	p_tmsi_sig	p_tmsi_sig	N/A

Table 30-42 GTP Information Elements (continued)

Value	Version 0	Version 1	Version 2
13	ms_validated	ms_validated	N/A
14	recovery	recovery	N/A
15	selection_mode	selection_mode	N/A
16	flow_label_data_1	teid_1	N/A
17	flow_label_signalling	teid_control	N/A
18	flow_label_data_2	teid_2	N/A
19	ms_unreachable	teardown_ind	N/A
20	N/A	nsapi	N/A
21	N/A	ranap	N/A
22	N/A	rab_context	N/A
23	N/A	radio_priority_sms	N/A
24	N/A	radio_priority	N/A
25	N/A	packet_flow_id	N/A
26	N/A	charging_char	N/A
27	N/A	trace_ref	N/A
28	N/A	trace_type	N/A
29	N/A	ms_unreachable	N/A
71	N/A	N/A	apn
72	N/A	N/A	ambr
73	N/A	N/A	ebi
74	N/A	N/A	ip_addr
75	N/A	N/A	mei
76	N/A	N/A	msisdn
77	N/A	N/A	indication
78	N/A	N/A	pco
79	N/A	N/A	paa
80	N/A	N/A	bearer_qos
80	N/A	N/A	flow_qos
82	N/A	N/A	rat_type
83	N/A	N/A	serving_network
84	N/A	N/A	bearer_tft
85	N/A	N/A	tad
86	N/A	N/A	uli
87	N/A	N/A	f_teid
88	N/A	N/A	tmsi
89	N/A	N/A	cn_id

Table 30-42 GTP Information Elements (continued)

Value	Version 0	Version 1	Version 2
90	N/A	N/A	s103pdf
91	N/A	N/A	s1udf
92	N/A	N/A	delay_value
93	N/A	N/A	bearer_context
94	N/A	N/A	charging_id
95	N/A	N/A	charging_char
96	N/A	N/A	trace_info
97	N/A	N/A	bearer_flag
99	N/A	N/A	pdn_type
100	N/A	N/A	pti
101	N/A	N/A	drx_parameter
103	N/A	N/A	gsm_key_tri
104	N/A	N/A	umts_key_cipher_quin
105	N/A	N/A	gsm_key_cipher_quin
106	N/A	N/A	umts_key_quin
107	N/A	N/A	eps_quad
108	N/A	N/A	umts_key_quad_quin
109	N/A	N/A	pdn_connection
110	N/A	N/A	pdn_number
111	N/A	N/A	p_tmsi
112	N/A	N/A	p_tmsi_sig
113	N/A	N/A	hop_counter
114	N/A	N/A	ue_time_zone
115	N/A	N/A	trace_ref
116	N/A	N/A	complete_request_msg
117	N/A	N/A	guti
118	N/A	N/A	f_container
119	N/A	N/A	f_cause
120	N/A	N/A	plmn_id
121	N/A	N/A	target_id
123	N/A	N/A	packet_flow_id
124	N/A	N/A	rab_contex
125	N/A	N/A	src_rnc_pdcp
126	N/A	N/A	udp_src_port
127	charge_id	charge_id	apn_restriction
128	end_user_address	end_user_address	selection_mode

Table 30-42 GTP Information Elements (continued)

Value	Version 0	Version 1	Version 2
129	mm_context	mm_context	src_id
130	pdp_context	pdp_context	N/A
131	apn	apn	change_report_action
132	protocol_config	protocol_config	fq_csid
133	gsn	gsn	channel
134	msisdn	msisdn	emlpp_pri
135	N/A	qos	node_type
136	N/A	authentication_qu	fqdn
137	N/A	tft	ti
138	N/A	target_id	mbms_session_duration
139	N/A	utran_trans	mbms_service_area
140	N/A	rab_setup	mbms_session_id
141	N/A	ext_header	mbms_flow_id
142	N/A	trigger_id	mbms_ip_multicast
143	N/A	omc_id	mbms_distribution_ack
144	N/A	ran_trans	rfsp_index
145	N/A	pdp_context_pri	uci
146	N/A	addi_rab_setup	csg_info
147	N/A	sgsn_number	csg_id
148	N/A	common_flag	cmi
149	N/A	apn_restriction	service_indicator
150	N/A	radio_priority_lcs	detach_type
151	N/A	rat_type	ldn
152	N/A	user_loc_info	node_feature
153	N/A	ms_time_zone	mbms_time_to_transfer
154	N/A	imei_sv	throttling
155	N/A	camel	arp
156	N/A	mbms_ue_context	epc_timer
157	N/A	tmp_mobile_group_id	signalling_priority_indication
158	N/A	rim_routing_addr	tmgi
159	N/A	mbms_config	mm_srvcc
160	N/A	mbms_service_area	flags_srvcc
161	N/A	src_rnc_pdcp	nmbr
162	N/A	addi_trace_info	N/A
163	N/A	hop_counter	N/A
164	N/A	plmn_id	N/A

Table 30-42 GTP Information Elements (continued)

Value	Version 0	Version 1	Version 2
165	N/A	mbms_session_id	N/A
166	N/A	mbms_2g3g_indicator	N/A
167	N/A	enhanced_nsapi	N/A
168	N/A	mbms_session_duration	N/A
169	N/A	addi_mbms_trace_info	N/A
170	N/A	mbms_session_repetition_num	N/A
171	N/A	mbms_time_to_data	N/A
173	N/A	bss	N/A
174	N/A	cell_id	N/A
175	N/A	pdu_num	N/A
177	N/A	mbms_bearer_capab	N/A
178	N/A	rim_routing_disc	N/A
179	N/A	list_pfc	N/A
180	N/A	ps_xid	N/A
181	N/A	ms_info_change_report	N/A
182	N/A	direct_tunnel_flags	N/A
183	N/A	correlation_id	N/A
184	N/A	bearer_control_mode	N/A
185	N/A	mbms_flow_id	N/A
186	N/A	mbms_ip_multicast	N/A
187	N/A	mbms_distribution_ack	N/A
188	N/A	reliable_inter_rat_handover	N/A
189	N/A	rfsp_index	N/A
190	N/A	fqdn	N/A
191	N/A	evolved_allocation1	N/A
192	N/A	evolved_allocation2	N/A
193	N/A	extended_flags	N/A
194	N/A	uci	N/A
195	N/A	csg_info	N/A
196	N/A	csg_id	N/A
197	N/A	cmi	N/A
198	N/A	apn_ambr	N/A
199	N/A	ue_network	N/A
200	N/A	ue_ambr	N/A
201	N/A	apn_ambr_nsapi	N/A
202	N/A	ggsn_backoff_timer	N/A

Table 30-42 GTP Information Elements (continued)

Value	Version 0	Version 1	Version 2
203	N/A	signalling_priority_indication	N/A
204	N/A	signalling_priority_indication_nsapi	N/A
205	N/A	high_bitrate	N/A
206	N/A	max_mbr	N/A
251	charging_gateway_addr	charging_gateway_addr	N/A
255	private_extension	private_extension	private_extension

You can use the following procedure to specify a GTP information element.

To specify a GTP information element:

Step 1 On the Create Rule page, select **gtp_info** in the drop-down list and click **Add Option.**

The `gtp_info` keyword appears.

Step 2 Specify a single defined decimal value 0 to 255 for the information element, or a single defined string. See the GTP Information Elements table for values and strings recognized by the system.

Modbus Keywords

License: Protection

You can use Modbus keywords to point to the beginning of the Data field in a Modbus request or response, to match against the Modbus Function Code, and to match against a Modbus Unit ID. You can use Modbus keywords alone or in combination with other keywords such as `content` and `byte_jump`.

See the following sections for more information:

- modbus_data, page 30-74
- modbus_func, page 30-75
- modbus_unit, page 30-75

modbus_data

You can use the `modbus_data` keyword to point to the beginning of the Data field in a Modbus request or response.

To point to the beginning of the Modbus Data field:

Step 1 On the Create Rule page, select **modbus_data** from the drop-down list and click **Add Option.**

The `modbus_data` keyword appears.

The `modbus_data` keyword has no arguments.

modbus_func

You can use the `modbus_func` keyword to match against the Function Code field in a Modbus application layer request or response header. You can specify either a single defined decimal value or a single defined string for a Modbus function code.

The following table lists the defined values and strings recognized by the system for Modbus function codes.

Table 30-43 Modbus Function Codes

Value	String
1	read_coils
2	read_discrete_inputs
3	read_holding_registers
4	read_input_registers
5	write_single_coil
6	write_single_register
7	read_exception_status
8	diagnostics
11	get_comm_event_counter
12	get_comm_event_log
15	write_multiple_coils
16	write_multiple_registers
17	report_slave_id
20	read_file_record
21	write_file_record
22	mask_write_register
23	read_write_multiple_registers
24	read_fifo_queue
43	encapsulated_interface_transport

To specify a Modbus function code:

Step 1 On the Create Rule page, select **modbus_func** in the drop-down list and click **Add Option.**

The `modbus_func` keyword appears.

Step 2 Specify a single defined decimal value 0 to 255 for the function code, or a single defined string. See the Modbus Function Codes table for values and strings recognized by the system.

modbus_unit

You can use the `modbus_unit` keyword to match a single decimal value against the Unit ID field in a Modbus request or response header.

To specify a Modbus unit ID:

Step 1 On the Create Rule page, select **modbus_unit** in the drop-down list and click **Add Option.**

The `modbus_unit` keyword appears.

Step 2 Specify a decimal value 0 through 255.

DNP3 Keywords

License: Protection

You can use DNP3 keywords to point to the beginning of application layer fragments, to match against DNP3 function codes and objects in DNP3 responses and requests, and to match against internal indication flags in DNP3 responses. You can use DNP3 keywords alone or in combination with other keywords such as `content` and `byte_jump`.

See the following sections for more information:

- dnp3_data, page 30-76
- dnp3_func, page 30-76
- dnp3_ind, page 30-78
- dnp3_obj, page 30-79

dnp3_data

You can use the `dnp3_data` keyword to point to the beginning of reassembled DNP3 application layer fragments.

The DNP3 preprocessor reassembles link layer frames into application layer fragments. The `dnp3_data` keyword points to the beginning of each application layer fragment; other rule options can match against the reassembled data within fragments without separating the data and adding checksums every 16 bytes.

To point to the beginning of reassembled DNP3 fragments:

Step 1 On the Create Rule page, select **modbus_data** from the drop-down list and click **Add Option.**

The `dnp3_data` keyword appears.

The `dnp3_data` keyword has no arguments.

dnp3_func

You can use the `dnp3_func` keyword to match against the Function Code field in a DNP3 application layer request or response header. You can specify either a single defined decimal value or a single defined string for a DNP3 function code.

The following table lists the defined values and strings recognized by the system for DNP3 function codes.

Table 30-44 DNP3 Function Codes

Value	String
0	confirm
1	read
2	write
3	select
4	operate
5	direct_operate
6	direct_operate_nr
7	immed_freeze
8	immed_freeze_nr
9	freeze_clear
10	freeze_clear_nr
11	freeze_at_time
12	freeze_at_time_nr
13	cold_restart
14	warm_restart
15	initialize_data
16	initialize_appl
17	start_appl
18	stop_appl
19	save_config
20	enable_unsolicited
21	disable_unsolicited
22	assign_class
23	delay_measure
24	record_current_time
25	open_file
26	close_file
27	delete_file
28	get_file_info
29	authenticate_file
30	abort_file
31	activate_config
32	authenticate_req
33	authenticate_err
129	response

Table 30-44 DNP3 Function Codes (continued)

Value	String
130	unsolicited_response
131	authenticate_resp

To specify DNP3 function codes:

Step 1 On the Create Rule page, select **dnp3_func** in the drop-down list and click **Add Option.**

The `dnp3_func` keyword appears.

Step 2 Specify a single defined decimal value 0 to 255 for the function code, or a single defined string. See the DNP3 Function Codes table for values and strings recognized by the system.

dnp3_ind

You can use the `dnp3_ind` keyword to match against flags in the Internal Indications field in a DNP3 application layer response header.

You can specify the string for a single known flag or a comma-separated list of flags, as seen in the following example:

```
class_1_events, class_2_events
```

When you specify multiple flags, the keyword matches against any flag in the list. To detect a combination of flags, use the `dnp3_ind` keyword multiple times in a rule.

The following list provides the string syntax recognized by the system for defined DNP3 internal indications flags.

```
class_1_events
class_2_events
class_3_events
need_time
local_control
device_trouble
device_restart
no_func_code_support
object_unknown
parameter_error
event_buffer_overflow
already_executing
config_corrupt
reserved_2
reserved_1
```

To specify DNP3 internal indications flags:

Step 1 On the Create Rule page, select **dnp3_ind** in the drop-down list and click **Add Option.**

The `dnp3_ind` keyword appears.

Step 2 You can specify the string for a single known flag or a comma-separated list of flags.

dnp3_obj

You can use the `dnp3_obj` keyword to match against DNP3 object headers in a request or response.

DNP3 data is comprised of a series of DNP3 objects of different types such as analog input, binary input, and so on. Each type is identified with a *group* such as analog input group, binary input group, and so on, each of which can be identified by a decimal value. The objects in each group are further identified by an *object variation* such as 16-bit integers, 32-bit integers, short floating point, and so on, each of which specifies the data format of the object. Each type of object variation can also be identified by a decimal value.

You identify object headers by specifying the decimal number for the type of object header group and the decimal number for the type of object variation. The combination of the two defines a specific type of DNP3 object.

To specify a DNP3 object:

Step 1 On the Create Rule page, select **dnp3_obj** in the drop-down list and click **Add Option.**

The `dnp3_obj` keyword appears.

Step 2 Specify a decimal value 0 through 255 to identify a known object group, and another decimal value 0 through 255 to identify a known object variation type.

Inspecting Packet Characteristics

License: Protection

You can write rules that only generate events against packets with specific packet characteristics. The ASA FirePOWER module provides the following keywords to evaluate packet characteristics:

- dsize, page 30-79
- isdataat, page 30-80
- sameip, page 30-81
- fragoffset, page 30-81
- cvs, page 30-81

dsize

License: Protection

The `dsize` keyword tests the packet payload size. With it, you can use the greater than and less than operators (`<` and `>`) to specify a range of values. You can use the following syntax to specify ranges:

```
>number_of_bytes
<number_of_bytes
number_of_bytes<>number_of_bytes
```

For example, to indicate a packet size greater than 400 bytes, use `>400` as the `dtype` value. To indicate a packet size of less than 500 bytes, use `<500`. To specify that the rule trigger against any packet between 400 and 500 bytes inclusive, use `400<>500`.

Caution The `dsize` keyword tests packets before they are decoded by any preprocessors.

isdataat

License: Protection

The `isdataat` keyword instructs the rules engine to verify that data resides at a specific location in the payload.

The following table lists the arguments you can use with the `isdataat` keyword.

Table 30-45 isdataat Arguments

Argument	Type	Description
Offset	Required	The specific location in the payload. For example, to test that data appears at byte 50 in the packet payload, you would specify `50` as the offset value. A `!` modifier negates the results of the `isdataat` test; it alerts if a certain amount of data is not present within the payload. You can also use an existing `byte_extract` variable to specify the value for this argument. See Reading Packet Data into Keyword Arguments, page 30-82 for more information.
Relative	Optional	Makes the location relative to the last successful content match. If you specify a relative location, note that the counter starts at byte 0, so calculate the location by subtracting 1 from the number of bytes you want to move forward from the last successful content match. For example, to specify that the data must appear at the ninth byte after the last successful content match, you would specify a relative offset of `8`.
Raw Data	Optional	Specifies that the data is located in the original packet payload before decoding or application layer normalization by any ASA FirePOWER module preprocessor. You can use this argument with **Relative** if the previous content match was in the raw packet data.

For example, in a rule searching for the content `foo`, if the value for `isdataat` is specified as the following:

- Offset = !10
- Relative = enabled

The system alerts if the rules engine does not detect 10 bytes after `foo` before the payload ends.

To use isdataat:

Step 1 On the Create Rule page, select `isdataat` in the drop-down list and click **Add Option**.

The `isdataat` section appears.

sameip

License: Protection

The `sameip` keyword tests that a packet's source and destination IP addresses are the same. It does not take an argument.

fragoffset

License: Protection

The `fragoffset` keyword tests the offset of a fragmented packet. This is useful because some exploits (such as WinNuke denial-of-service attacks) use hand-generated packet fragments that have specific offsets.

For example, to test whether the offset of a fragmented packet is 31337 bytes, specify `31337` as the `fragoffset` value.

You can use the following operators when specifying arguments for the `fragoffset` keyword.

Table 30-46 fragoffset Keyword Argument Operators

Operator	Description
!	not
>	greater than
<	less than

Note that you cannot use the not (`!`) operator in combination with `<` or `>`.

cvs

License: Protection

The `cvs` keyword tests Concurrent Versions System (CVS) traffic for malformed CVS entries. An attacker can use a malformed entry to force a heap overflow and execute malicious code on the CVS server. This keyword can be used to identify attacks against two known CVS vulnerabilities: CVE-2004-0396 (CVS 1.11.x up to 1.11.15, and 1.12.x up to 1.12.7) and CVS-2004-0414 (CVS 1.12.x through 1.12.8, and 1.11.x through 1.11.16). The `cvs` keyword checks for a well-formed entry, and generates alerts when a malformed entry is detected.

Your rule should include the ports where CVS runs. In addition, any ports where traffic may occur should be added to the list of ports for stream reassembly in your TCP policies so state can be maintained for CVS sessions. The TCP ports 2401 (`pserver`) and 514 (`rsh`) are included in the list of client ports where stream reassembly occurs. However, note that if your server runs as an `xinetd` server (i.e., pserver), it can run on any TCP port. Add any non-standard ports to the stream reassembly **Client Ports** list. For more information, see Selecting Stream Reassembly Options, page 24-26.

To detect malformed CVS entries:

Step 1 Add the `cvs` option to a rule and type `invalid-entry` as the keyword argument.

Reading Packet Data into Keyword Arguments

License: Protection

You can use the `byte_extract` keyword to read a specified number of bytes from a packet into a variable. You can then use the variable later in the same rule as the value for specific arguments in certain other detection keywords.

This is useful, for example, for extracting data size from packets where a specific segment of bytes describes the number of bytes included in data within the packet. For example, a specific segment of bytes might say that subsequent data is comprised of four bytes; you can extract the data size of four bytes to use as your variable value.

You can use `byte_extract` to create up to two separate variables in a rule concurrently. You can redefine a `byte_extract` variable any number of times; entering a new `byte_extract` keyword with the same variable name and a different variable definition overwrites the previous definition of that variable.

The following table describes the arguments required by the `byte_extract` keyword.

Table 30-47 Required byte_extract Arguments

Argument	Description
Bytes to Extract	The number of bytes to extract from the packet. You can specify 1, 2, 3, or 4 bytes.
Offset	The number of bytes into the payload to begin extracting data. You can specify -65534 to 65535 bytes. The offset counter starts at byte 0, so calculate the offset value by subtracting 1 from the number of bytes you want to count forward. For example, specify `7` to count forward 8 bytes. The rules engine counts forward from the beginning of the packet payload or, if you also specify **Relative**, after the last successful content match. Note that you can specify negative numbers only when you also specify **Relative**; see the Additional Optional byte_extract Arguments table for more information.
Variable Name	The variable name to use in arguments for other detection keywords. You can specify an alphanumeric string that must begin with a letter.

To further define how the system locates the data to extract, you can use the arguments described in the following table.

Table 30-48 Additional Optional byte_extract Arguments

Argument	Description
Multiplier	A multiplier for the value extracted from the packet. You can specify 0 to 65535. If you do not specify a multiplier, the default value is 1.
Align	Rounds the extracted value to the nearest 2-byte or 4-byte boundary. When you also select **Multiplier**, the system applies the multiplier before the alignment.
Relative	Makes **Offset** relative to the end of the last successful content match instead of the beginning of the payload. See the Required byte_extract Arguments table for more information.

You can specify only one of **DCE/RPC**, **Endian**, or **Number Type**.

To define how the `byte_extract` keyword calculates the bytes it tests, you can choose from the arguments in the following table. The rules engine uses big endian byte order if you do not select either argument.

Table 30-49 Endianness byte_extract Arguments

Argument	Description
Big Endian	Processes data in big endian byte order, which is the default network byte order.
Little Endian	Processes data in little endian byte order.
DCE/RPC	Specifies a `byte_extract` keyword for traffic processed by the DCE/RPC preprocessor. See Decoding DCE/RPC Traffic, page 22-2 for more information. The DCE/RPC preprocessor determines big endian or little endian byte order, and the **Number Type** and **Endian** arguments do not apply. When you enable this argument, you can also use `byte_extract` in conjunction with other specific DCE/RPC keywords. See DCE/RPC Keywords, page 30-59 for more information.

You can specify a number type to read data as an ASCII string. To define how the system views string data in a packet, you can select one of the arguments in the following table.

Table 30-50 Number Type byte_extract arguments

Argument	Description
Hexadecimal String	Reads extracted string data in hexadecimal format.
Decimal String	Reads extracted string data in decimal format.
Octal String	Reads extracted string data in octal format.

For example, if the value for `byte_extract` is specified as the following:

- Bytes to Extract = 4
- Variable Name = var
- Offset = 8
- Relative = enabled

the rules engine reads the number described in the four bytes that appear 9 bytes away from (relative to) the last successful content match into a variable named `var`, which you can specify later in the rule as the value for certain keyword arguments.

The following table lists the keyword arguments where you can specify a variable defined in the `byte_extract` keyword.

Table 30-51 Arguments Accepting a byte_extract Variable

Keyword	Argument	For more information, see...
content	Depth, Offset, Distance, Within	Constraining Content Matches, page 30-17
byte_jump	Offset	byte_jump, page 30-31
byte_test	Offset, Value	byte_test, page 30-33
isdataat	Offset	isdataat, page 30-80

To use byte_extract:

Step 1 On the Create Rule page, select t `byte_extract` in the drop-down list and click **Add Option**.

The `byte_extract` section appears beneath the last keyword you selected.

Initiating Active Responses with Rule Keywords

License: Protection

The system can initiate active responses to close TCP connections in response to triggered TCP rules or UDP sessions in response to triggered UDP rules. Two keywords provide you with separate approaches to initiating active responses. When a packet triggers a rule containing either of the keywords, the system initiates a single active response. You can also use the `config response` command to configure the active response interface to use and the number of TCP resets to attempt in a passive deployment.

Active responses are most effective in inline deployments because resets are more likely to arrive in time to affect the connection or session. For example, in response to the `react` keyword in an inline deployment, the system inserts a TCP reset (RST) packet directly into the traffic for each end of the connection, which normally should close the connection.

Active responses are not intended to take the place of a firewall for a number of reasons, including that the system cannot insert packets in passive deployments and an attacker may have chosen to ignore or circumvent active responses.

Because active responses can be routed back, the system does not allow TCP resets to initiate TCP resets; this prevents an unending sequence of active responses. The system also does not allow ICMP unreachable packets to initiate ICMP unreachable packets in keeping with standard practice.

You can configure the TCP stream preprocessor to detect additional traffic on a connection or session after an intrusion rule has triggered an active response. When the preprocessor detects additional traffic, it sends additional active responses up to a specified maximum to both ends of the connection or session. See Initiating Active Responses with Intrusion Drop Rules, page 24-2 for more information.

See the following sections for information specific to the keywords you can use to initiate active responses:

- Initiating Active Responses by Type and Direction, page 30-84
- Sending an HTML Page Before a TCP Reset, page 30-85
- Setting the Active Response Reset Attempts and Interface, page 30-86

Initiating Active Responses by Type and Direction

License: Protection

You can use the `resp` keyword to actively respond to TCP connections or UDP sessions, depending on whether you specify the TCP or UDP protocol in the rule header. See Specifying Protocols, page 30-4 for more information.

Keyword arguments allow you to specify the packet direction and whether to use TCP reset (RST) packets or ICMP unreachable packets as active responses.

You can use any of the TCP reset or ICMP unreachable arguments to close TCP connections. You should use only ICMP unreachable arguments to close UDP sessions.

Different TCP reset arguments also allow you to target active responses to the packet source, destination, or both. All ICMP unreachable arguments target the packet source and allow you to specify whether to use an ICMP network, host, or port unreachable packet, or all three.

The following table lists the arguments you can use with the `resp` keyword to specify exactly what you want the ASA FirePOWER module to do when the rule triggers.

Table 30-52 resp Arguments

Argument	Description
reset_source	Directs a TCP reset packet to the endpoint that sent the packet that triggered the rule. Alternatively, you can specify `rst_snd`, which is supported for backward compatibility.
reset_dest	Directs a TCP reset packet to the intended destination endpoint of the packet that triggered the rule. Alternatively, you can specify `rst_rcv`, which is supported for backward compatibility.
reset_both	Directs a TCP reset packet to both the sending and receiving endpoints. Alternatively, you can specify `rst_all`, which is supported for backward compatibility.
icmp_net	Directs an ICMP network unreachable message to the sender.
icmp_host	Directs an ICMP host unreachable message to the sender.
icmp_port	Directs an ICMP port unreachable message to the sender. This argument is used to terminate UDP traffic.
icmp_all	Directs the following ICMP messages to the sender: • network unreachable • host unreachable • port unreachable

For example, to configure a rule to reset both sides of a connection when a rule is triggered, use `reset_both` as the value for the `resp` keyword.

You can use a comma-separated list to specify multiple arguments as follows:

`argument,argument,argument`

See Setting the Active Response Reset Attempts and Interface, page 30-86 for information on using the `config response` command to configure the active response interface to use and the number of TCP resets to attempt in a passive deployment.

To specify active responses:

Step 1 On the Create Rule page, select **resp** in the drop-down list and click **Add Option.**

The `resp` keyword appears.

Step 2 Specify any of the arguments in the resp Arguments table in the **resp** field; use a comma-separated list to specify multiple arguments.

Sending an HTML Page Before a TCP Reset

License: Protection

You can use the `react` keyword to send a default HTML page to the TCP connection client when a packet triggers the rule; after sending the HTML page, the system uses TCP reset packets to initiate active responses to both ends of the connection. The `react` keyword does not trigger active responses for UDP traffic.

Optionally, you can specify the following argument:

`msg`

When a packet triggers a `react` rule that uses the `msg` argument, the HTML page includes the rule event message. See Understanding Rule Anatomy, page 30-2 for a description of the event message field.

If you do not specify the `msg` argument, the HTML page includes the following message:

```
You are attempting to access a forbidden site.
Consult your system administrator for details.
```

Note Because active responses can be routed back, ensure that the HTML response page does not trigger a `react` rule; this could result in an unending sequence of active responses. Cisco recommends that you test `react` rules extensively before activating them in a production environment.

See Setting the Active Response Reset Attempts and Interface, page 30-86 for information on using the `config response` command to configure the active response interface to use and the number of TCP resets to attempt in a passive deployment.

To send an HTML page before initiating an active responses:

Step 1 On the Create Rule page, select **react** in the drop-down list and click **Add Option.**

The `react` keyword appears.

Step 2 You have two choices:

- To send an HTML page that includes the event message configured for the rule to the client before closing a connection, type `msg` in the **react** field.
- To send an HTML page that includes the following default message to the client before closing a connection, leave the **react** field blank:

```
You are attempting to access a forbidden site.
Consult your system administrator for details
```

Setting the Active Response Reset Attempts and Interface

License: Protection

You can use the **config response** command to further configure the behavior of TCP resets initiated by `resp` and `react` rules. This command also affects the behavior of active responses initiated by drop rules; see Initiating Active Responses with Intrusion Drop Rules, page 24-2 for more information.

You use the **config response** command by inserting it on a separate line in the USER_CONF advanced variable. See Understanding Advanced Variables, page 2-27 for information on using a USER_CONF variable.

Caution Do **not** use the `USER_CONF` advanced variable to configure an intrusion policy feature unless you are instructed to do so in the feature description or by Support. Conflicting or duplicate configurations will halt the system.

To specify active response reset attempts, the active response interface, or both:

Step 1 Depending on whether you want to specify only the number of active responses, only the active response interface, or both, insert a form of the `config response` command on a separate line in the USER_CONF advanced variable. You have the following choices:

- To specify only the number of active response attempts, insert the command:
  `config response: attempts att`

  For example: `config response: attempts 10`
- To specify only the active response interface, insert the command:
  `config response: device dev`

  For example: `config response: device eth0`
- To specify both the number of active response attempts and the active response interface, insert the command:
  `config response: attempts att, device dev`

  For example: `config response: attempts 10, device eth0`

  where:

  att is the number 1 to 20 of attempts to land each TCP reset packet within the current connection window so the receiving host accepts the packet. This sequence *strafing* is useful only in passive deployments; in inline deployments, the system inserts reset packets directly into the stream in place of triggering packets. the system sends only 1 ICMP reachable active response.

  dev is an alternate interface where you want the system to send active responses in a passive deployment or insert active responses in an inline deployment.

Filtering Events

License: Protection

You can use the `detection_filter` keyword to prevent a rule from generating events unless a specified number of packets trigger the rule within a specified time. This can stop the rule from prematurely generating events. For example, two or three failed login attempts within a few seconds could be expected behavior, but a large number of attempts within the same time could indicate a brute force attack.

The `detection_filter` keyword requires arguments that define whether the system tracks the source or destination IP address, the number of times the detection criteria must be met before triggering an event, and how long to continue the count.

Use the following syntax to delay the triggering of events:

`track by_src/by_dst, count count, seconds number_of_seconds`

The `track` argument specifies whether to use the packet's source or destination IP address when counting the number of packets that meet the rule's detection criteria. Select from the argument values described in the following table to specify how the system tracks event instances.

Table 30-53 detection_filter Track Arguments

Argument	Description
by_src	Detection criteria count by source IP address.
by_dst	Detection criteria count by destination IP address.

The `count` argument specifies the number of packets that must trigger the rule for the specified IP address within the specified time before the rule generates an event.

The `seconds` argument specifies the number of seconds within which the specified number of packets must trigger the rule before the rule generates an event.

Consider the case of a rule that searches packets for the content `foo` and uses the `detection_filter` keyword with the following arguments:

```
track by_src, count 10, seconds 20
```

In the example, the rule will not generate an event until it has detected `foo` in 10 packets within 20 seconds from a given source IP address. If the system detects only 7 packets containing `foo` within the first 20 seconds, no event is generated. However, if `foo` occurs 40 times in the first 20 seconds, the rule generates 30 events and the count begins again when 20 seconds have elapsed.

Comparing the threshold and detection_filter Keywords

The `detection_filter` keyword replaces the deprecated `threshold` keyword. The `threshold` keyword is still supported for backward compatibility and operates the same as thresholds that you set within an intrusion policy.

The `detection_filter` keyword is a detection feature that is applied before a packet triggers a rule. The rule does not generate an event for triggering packets detected before the specified packet count and, in an inline deployment, does not drop those packets if the rule is set to drop packets. Conversely, the rule does generate events for packets that trigger the rule and occur after the specified packet count and, in an inline deployment, drops those packets if the rule is set to drop packets.

Thresholding is an event notification feature that does not result in a detection action. It is applied after a packet triggers an event. In an inline deployment, a rule that is set to drop packets drops all packets that trigger the rule, independent of the rule threshold.

Note that you can use the `detection_filter` keyword in any combination with the intrusion event thresholding, intrusion event suppression, and rate-based attack prevention features in an intrusion policy. Note also that policy validation fails if you enable an imported local rule that uses the deprecated `threshold` keyword in combination with the intrusion event thresholding feature in an intrusion policy. See Configuring Event Thresholding, page 27-21, Configuring Suppression Per Intrusion Policy, page 27-25, Setting a Dynamic Rule State, page 27-29, and Importing Local Rule Files, page 46-14 for more information.

Evaluating Post-Attack Traffic

License: Protection

Use the `tag` keyword to tell the system to log additional traffic for the host or session. Use the following syntax when specifying the type and amount of traffic you want to capture using the `tag` keyword:

```
tagging_type, count, metric, optional_direction
```

The next three tables describe the other available arguments.

You can choose from two types of tagging. The following table describes the two types of tagging. Note that the session tag argument type causes the system to log packets from the same session as if they came from different sessions if you configure only rule header options in the intrusion rule. To group packets from the same session together, configure one or more rule options (such as a `flag` keyword or `content` keyword) within the same intrusion rule.

Table 30-54 Tag Arguments

Argument	Description
session	Logs packets in the session that triggered the rule.
host	Logs packets from the host that sent the packet that triggered the rule. You can add a directional modifier to log only the traffic coming from the host (`src`) or going to the host (`dst`).

To indicate how much traffic you want to log, use the following argument:

Table 30-55 Count Argument

Argument	Description
count	The number of packets or seconds you want to log after the rule triggers. This unit of measure is specified with the metric argument, which follows the count argument.

Select the metric you want to use to log by time or volume of traffic from those described in the following table.

Caution High-bandwidth networks can see thousands of packets per second, and tagging a large number of packets may seriously affect performance, so make sure you tune this setting for your network environment.

Table 30-56 Logging Metrics Arguments

Argument	Description
packets	Logs the number of packets specified by the count after the rule triggers.
seconds	Logs traffic for the number of seconds specified by the count after the rule triggers.

For example, when a rule with the following `tag` keyword value triggers:

```
host, 30, seconds, dst
```

all packets that are transmitted from the client to the host for the next 30 seconds are logged.

Detecting Attacks That Span Multiple Packets

License: Protection

Use the `flowbits` keyword to assign state names to sessions. By analyzing subsequent packets in a session according to the previously named state, the system can detect and alert on exploits that span multiple packets in a single session.

The `flowbits` state name is a user-defined label assigned to packets in a specific part of a session. You can label packets with state names based on packet content to help distinguish malicious packets from those you do not want to alert on. You can define up to 1024 state names. For example, if you want to alert on malicious packets that you know only occur after a successful login, you can use the `flowbits` keyword to filter out the packets that constitute an initial login attempt so you can focus only on the malicious packets. You can do this by first creating a rule that labels all packets in the session that have an established login with a `logged_in` state, then creating a second rule where `flowbits` checks for packets with the state you set in the first rule and acts only on those packets. See flowbits Example Using state_name, page 30-91 for an example that uses `flowbits` to determine if a user is logged in.

An optional *group name* allows you to include a state name in a group of states. A state name can belong to several groups. States not associated with a group are not mutually exclusive, so a rule that triggers and sets a state that is not associated with a group does not affect other currently set states. See flowbits Example Resulting in a False Positive, page 30-92 for an example that illustrates how including a state name in a group can prevent false positives by unsetting another state in the same group.

The following table describes the various combinations of operators, states, and groups available to the `flowbits` keyword. Note that state names can contain alphanumeric characters, periods (.), underscores (_), and dashes (-).

Table 30-57 flowbits Options

Operator	State Option	Group	Description
`set`	`state_name`	optional	Sets the specified state for a packet. Sets the state in the specified group if a group is defined.
	`state_name&state_name`	optional	Sets the specified states for a packet. Sets the states in the specified group if a group is defined.
`setx`	`state_name`	mandatory	Sets the specified state in the specified group for a packet, and unsets all other states in the group.
	`state_name&state_name`	mandatory	Sets the specified states in the specified group for a packet, and unsets all other states in the group.
`unset`	`state_name`	no group	Unsets the specified state for a packet.
	`state_name&state_name`	no group	Unsets the specified states for a packet.
	`all`	mandatory	Unsets all the states in the specified group.
`toggle`	`state_name`	no group	Unsets the specified state if it is set, and sets the specified state if it is unset.
	`state_name&state_name`	no group	Unsets the specified states if they are set, and sets the specified states if they are unset.
	`all`	mandatory	Unsets all states set in the specified group, and sets all states unset in the specified group.
`isset`	`state_name`	no group	Determines if the specified state is set in the packet.
	`state_name&state_name`	no group	Determines if the specified states are set in the packet.
	`state_name|state_name`	no group	Determines if any of the specified states are set in the packet.
	`any`	mandatory	Determines if any state is set in the specified group.
	`all`	mandatory	Determines if all states are set in the specified group.

Table 30-57 flowbits Options (continued)

Operator	State Option	Group	Description
isnotset	state_name	no group	Determines if the specified state is not set in the packet.
	state_name&state_name	no group	Determines if the specified states are not set in the packet.
	state_name|state_name	no group	Determines if any of the specified states is not set in the packet.
	any	mandatory	Determines if any state is not set in the packet.
	all	mandatory	Determines if all states are not set in the packet.
reset	(no state)	optional	Unsets all states for all packets. Unsets all states in a group if a group is specified.
noalert	(no state)	no group	Use this in conjunction with any other operator to suppress event generation.

Note the following when using the `flowbits` keyword:

- When using the `setx` operator, the specified state can only belong to the specified group, and not to any other group.
- You can define the `setx` operator multiple times, specifying different states and the same group with each instance.
- When you use the `setx` operator and specify a group, you cannot use the `set`, `toggle`, or `unset` operators on that specified group.
- The `isset` and `isnotset` operators evaluate for the specified state regardless of whether the state is in a group.
- During intrusion policy saves, intrusion policy reapplies, and access control policy applies (regardless of whether the access control policy references one intrusion policy or multiple intrusion policies), if you enable a rule that contains the `isset` or `isnotset` operator **without** a specified group, and you do not enable at least one rule that affects `flowbits` assignment (`set`, `setx`, `unset`, `toggle`) for the corresponding state name and protocol, all rules that affect `flowbits` assignment for the corresponding state name are enabled.
- During intrusion policy saves, intrusion policy reapplies, and access control policy applies (regardless of whether the access control policy references one intrusion policy or multiple intrusion policies), if you enable a rule that contains the `isset` or `isnotset` operator **with** a specified group, all rules that affect `flowbits` assignment (`set`, `setx`, `unset`, `toggle`) and define a corresponding group name are also enabled.

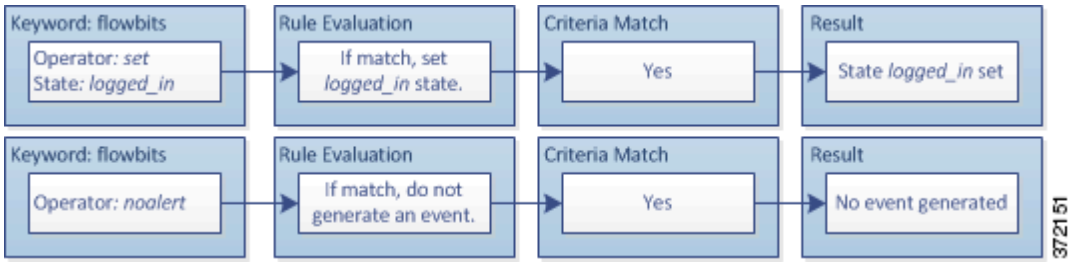
flowbits Example Using state_name

Consider the IMAP vulnerability described in Bugtraq ID #1110. This vulnerability exists in an implementation of IMAP, specifically in the LIST, LSUB, RENAME, FIND, and COPY commands. However, to take advantage of the vulnerability, the attacker must be logged into the IMAP server. Because the LOGIN confirmation from the IMAP server and the exploit that follows are necessarily in different packets, it is difficult to construct non-flow-based rules that catch this exploit. Using the `flowbits` keyword, you can construct a series of rules that track whether the user is logged into the IMAP server and, if so, generate an event if one of the attacks is detected. If the user is not logged in, the attack cannot exploit the vulnerability and no event is generated.

The two rule fragments that follow illustrate this example. The first rule fragment looks for an IMAP login confirmation from the IMAP server:

```
alert tcp any 143 -> any any (msg:"IMAP login"; content:"OK
LOGIN"; flowbits:set,logged_in; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

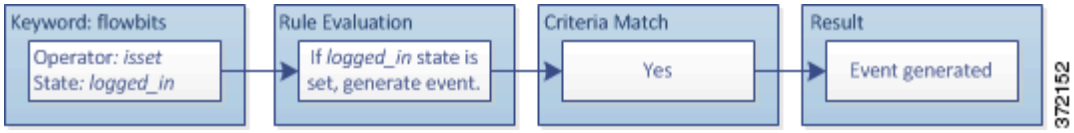


Note that `flowbits:set` sets a state of `logged_in`, while `flowbits:noalert` suppresses the alert because you are likely to see many innocuous login sessions on an IMAP server.

The next rule fragment looks for a LIST string, but does not generate an event unless the logged_in state has been set as a result of some previous packet in the session:

```
alert tcp any any -> any 143 (msg:"IMAP LIST";
content:"LIST"; flowbits:isset,logged_in;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:



In this case, if a previous packet has caused a rule containing the first fragment to trigger, then a rule containing the second fragment triggers and generates an event.

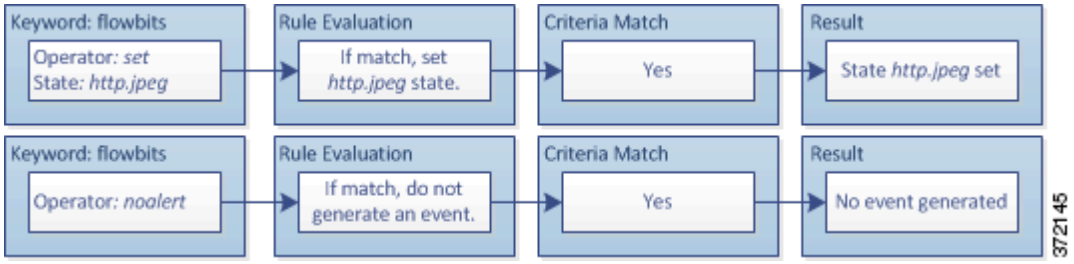
flowbits Example Resulting in a False Positive

Including different state names that are set in different rules in a group can prevent false positive events that might otherwise occur when content in a subsequent packet matches a rule whose state is no longer valid. The following example illustrates how you can get false positives when you do not include multiple state names in a group.

Consider the case where the following three rule fragments trigger in the order shown during a single session:

```
(msg:"JPEG transfer"; content:"image/";pcre:"/^Content-
Type\x3a(\s*|\s*\r?\n\s+)image\x2fp?jpe?g/smi";
flowbits:set,http.jpeg; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

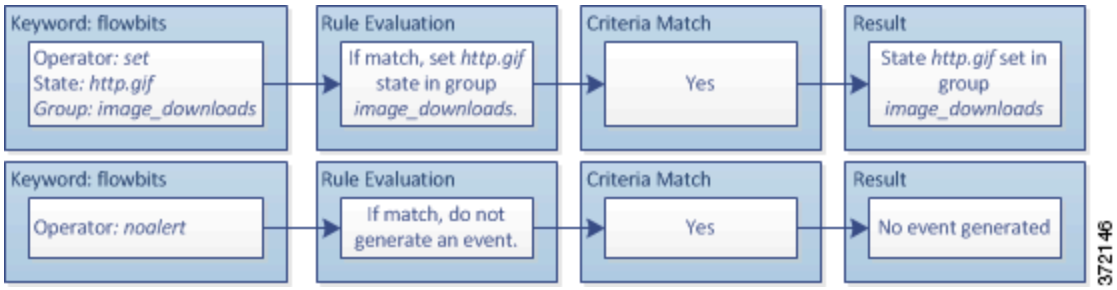


The `content` and `pcre` keywords in the first rule fragment match a JPEG file download, `flowbits:set,http.jpeg` sets the `http.jpeg` `flowbits` state, and `flowbits:noalert` stops the rule from generating events. No event is generated because the rule's purpose is to detect the file download and set the `flowbits` state so one or more companion rules can test for the state name in combination with malicious content and generate events when malicious content is detected.

The next rule fragment detects a GIF file download subsequent to the JPEG file download above:

```
(msg:"GIF transfer"; content:"image/"; pcre:"/^Content-
Type\x3a(\s*|\s*\r?\n\s+)image\x2fgif/smi";
flowbits:set,http.tif,image_downloads; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

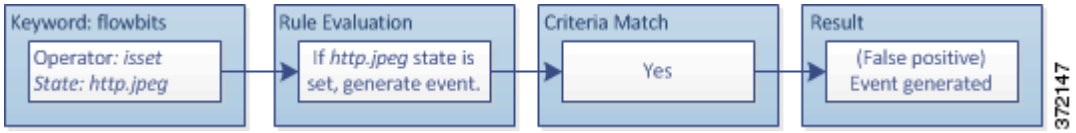


The `content` and `pcre` keywords in the second rule match the GIF file download, `flowbits:set,http.tif` sets the `http.tif` flowbit state, and `flowbits:noalert` stops the rule from generating an event. Note that the `http.jpeg` state set by the first rule fragment is still set even though it is no longer needed; this is because the JPEG download must have ended if a subsequent GIF download has been detected.

The third rule fragment is a companion to the first rule fragment:

```
(msg:"JPEG exploit";
flowbits:isset,http.jpeg;content:"|FF|"; pcre:"
/\xFF[\xE1\xE2\xED\xFE]\x00[\x00\x01]/";)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:



In the third rule fragment, `flowbits:isset,http.jpeg` determines that the now-irrelevant `http.jpeg` state is set, and `content` and `pcre` match content that would be malicious in a JPEG file but not in a GIF file. The third rule fragment results in a false positive event for a nonexistent exploit in a JPEG file.

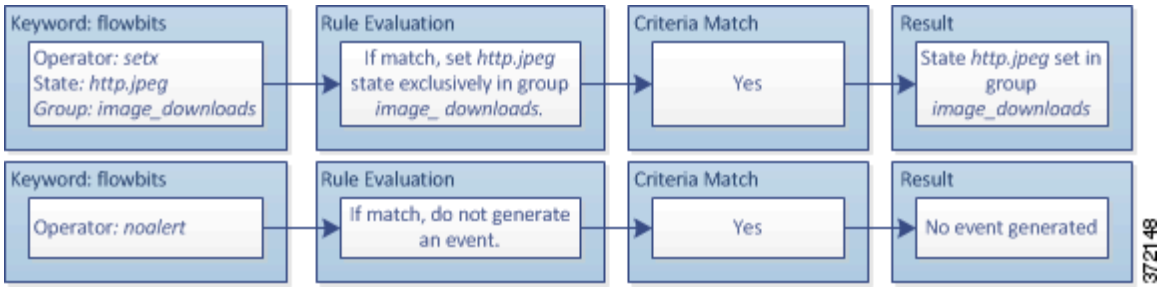
flowbits Example for Preventing False Positives

The following example illustrates how including state names in a group and using the `setx` operator can prevent false positives.

Consider the same case as the previous example, except that the first two rules now include their two different state names in the same state group.

```
(msg:"JPEG transfer"; content:"image/";pcre:"/^Content-
Type\x3a(\s*|\s*\r?\n\s+)image\x2fp?jpe?g/smi";
flowbits:setx,http.jpeg,image_downloads; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

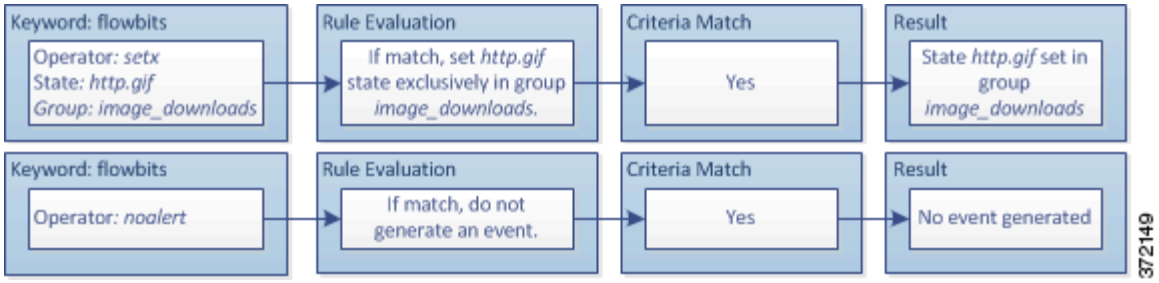


When the first rule fragment detects a JPEG file download, the `flowbits:setx,http.jpeg,image_downloads` keyword sets the `flowbits` state to `http.jpeg` and includes the state in the `image_downloads` group.

The next rule then detects a subsequent GIF file download:

```
(msg:"GIF transfer"; content:"image/"; pcre:"/^Content-
Type\x3a(\s*|\s*\r?\n\s+)image\x2fgif/smi";
flowbits:setx,http.tif,image_downloads; flowbits:noalert;)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:

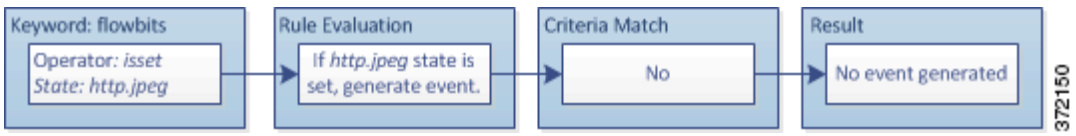


When the second rule fragment matches the GIF download, the `flowbits:setx,http.tif,image_downloads` keyword sets the `http.tif flowbits` state and unsets `http.jpeg`, the other state in the group.

The third rule fragment does not result in a false positive:

```
(msg:"JPEG exploit";
flowbits:isset,http.jpeg;content:"|FF|"; pcre:"/
\xFF[\xE1\xE2\xED\xFE]\x00[\x00\x01]/";)
```

The following diagram illustrates the effect of the `flowbits` keyword in the preceding rule fragment:



Because `flowbits:isset,http.jpeg` is false, the rules engine stops processing the rule and no event is generated, thus avoiding a false positive even in a case where content in the GIF file matches exploit content for a JPEG file.

Generating Events on the HTTP Encoding Type and Location

License: Protection

You can use the `http_encode` keyword to generate events on the type of encoding in an HTTP request or response before normalization, either in the HTTP URI, in non-cookie data in an HTTP header, in cookies in HTTP requests headers, or set-cookie data in HTTP responses.

You must configure the HTTP Inspect preprocessor to inspect HTTP responses and HTTP cookies to return matches for rules using the `http_encode` keyword. See Decoding HTTP Traffic, page 22-31 and Selecting Server-Level HTTP Normalization Options, page 22-33 for more information.

Also, you must enable both the decoding and alerting option for each specific encoding type in your HTTP Inspect preprocessor configuration for the `http_encode` keyword in an intrusion rule to trigger events on that encoding type. See Selecting Server-Level HTTP Normalization Encoding Options, page 22-41 for more information.

Note that the base36 encoding type has been deprecated. For backward compatibility, the base36 argument is allowed in existing rules, but it does not cause the rules engine to inspect base36 traffic.

The following table describes the encoding types this option can generate events for in HTTP URIs, headers, cookies, and set-cookies:

Table 30-58 http_encode Encoding Types

Encoding Type	Description
utf8	Detects UTF-8 encoding in the specified location when this encoding type is enabled for decoding by the HTTP Inspect preprocessor.
double_encode	Detects double encoding in the specified location when this encoding type is enabled for decoding by the HTTP Inspect preprocessor.
non_ascii	Detects non-ASCII characters in the specified location when non-ASCII characters are detected but the detected encoding type is not enabled.
uencode	Detects Microsoft %u encoding in the specified location when this encoding type is enabled for decoding by the HTTP Inspect preprocessor.
bare_byte	Detects bare byte encoding in the specified location when this encoding type is enabled for decoding by the HTTP Inspect preprocessor.

To identify the HTTP encoding type and location in an intrusion rule:

Step 1 Add the `http_encode` keyword to a rule.

Step 2 From the **Encoding Location** drop-down list, select whether to search for the specified encoding type in an HTTP URI, header, or cookie, including a set-cookie.

Step 3 Specify one or more encoding types using one of the following formats:

```
encode_type
encode_type|encode_type|encode_type...
!encode_type
```

where *encode_type* is one of the following:

```
utf8, double_encode, non_ascii, uencode, bare_byte
```

Note that you cannot use the negation (`!`) and OR (`|`) operators together.

Step 4 Optionally, add multiple `http_encode` keywords to the same rule to AND the conditions for each. For example, enter two keywords with the following conditions:

First `http_encode` keyword:

- **Encoding Location**: `HTTP URI`

- **Encoding Type**: `utf8`

Additional `http_encode` keyword:

- **Encoding Location**: `HTTP URI`
- **Encoding Type**: `uencode`

The example configuration searches the HTTP URI for UTF-8 AND Microsoft IIS %u encoding.

Detecting File Types and Versions

License: Protection

The `file_type` and `file_group` keywords allow you to detect files transmitted via FTP, HTTP, SMTP, IMAP, POP3, and NetBIOS-ssn (SMB) based on their type and version. Do **not** use more than one `file_type` or `file_group` keyword in a single intrusion rule.

Tip Updating your vulnerability database (VDB) populates the rule editor with the most up-to-date file types, versions, and groups. For more information, see Updating the Vulnerability Database, page 46-8.

You **must** enable specific preprocessors in order to generate intrusion events for traffic matching your `file_type` or `file_group` keywords.

Table 30-59 file_type and file_group Intrusion Event Generation

Transmission Protocol	Required Preprocessor or Preprocessor Option
FTP	FTP/Telnet preprocessor and the **Normalize TCP Payload** inline normalization preprocessor option; see Decoding FTP and Telnet Traffic, page 22-18 and Normalizing Inline Traffic, page 24-6.
HTTP	HTTP Inspect preprocessor; see Decoding HTTP Traffic, page 22-31.
SMTP	SMTP preprocessor; see Decoding SMTP Traffic, page 22-60.
IMAP	IMAP preprocessor; see Decoding IMAP Traffic, page 22-54.
POP3	POP preprocessor; see Decoding POP Traffic, page 22-57.
NetBIOS-ssn (SMB)	the **SMB File Inspection** DCE/RPC preprocessor option; see Decoding DCE/RPC Traffic, page 22-2.

For more information, see the following sections:

- file_type, page 30-96
- file_group, page 30-97

file_type

The `file_type` keyword allows you to specify the file type and version of a file detected in traffic. File type arguments (for example, **JPEG** and **PDF**) identify the format of the file you want to find in traffic.

Note Do **not** use the `file_type` keyword with another `file_type` or `file_group` keyword in the same intrusion rule.

The system selects **Any Version** by default, but some file types allow you to select version options (for example, PDF version **1.7**) to identify specific file type versions you want to find in traffic.

To view and configure the most up-to-date file types and versions, update your VDB. For more information, see Updating the Vulnerability Database, page 46-8.

To select file types and versions in an intrusion rule:

Step 1 On the Create Rule page, select **file_type** from the drop-down list and click **Add Option.**

The `file_type` keyword appears.

Step 2 Select one or more file types from the drop-down list. Selecting a file type automatically adds the argument to the rule.

To remove a file type argument from the rule, click the delete icon ( ) next to the file type you want to remove.

Step 3 Optionally, customize the target versions for each file type. The system selects **Any Version** by default, but some file types allow you to select individual target versions.

Note Updating your VDB populates the rule editor with the most up-to-date file types and versions. If you select **Any Version**, the system configures your rule to include new versions when they are added in later VDB updates.

file_group

The `file_group` keyword allows you to select a Cisco-defined group of similar file types to find in traffic (for example, **multimedia** or **audio**). File groups also include Cisco-defined versions for each file type in the group.

Note Do **not** use the `file_group` keyword with another `file_group` or `file_type` keyword in the same intrusion rule.

To view and configure the most up-to-date file groups, update your VDB. For more information, see Updating the Vulnerability Database, page 46-8.

To select a file group in an intrusion rule:

Step 1 On the Create Rule page, select **file_group** from the drop-down list and click **Add Option.**

The `file_group` keyword appears.

Step 2 Select a file group to add to the rule.

Pointing to a Specific Payload Type

License: Protection

The `file_data` keyword provides a pointer that serves as a reference for the positional arguments available for other keywords such as `content`, `byte_jump`, `byte_test`, and `pcre`. The detected traffic determines the type of data the `file_data` keyword points to. You can use the `file_data` keyword to point to the beginning of the following payload types:

- HTTP response body

  To inspect HTTP response packets, the HTTP Inspect preprocessor must be enabled and you must configure the preprocessor to inspect HTTP responses. See Decoding HTTP Traffic, page 22-31 and **Inspect HTTP Responses** in Selecting Server-Level HTTP Normalization Options, page 22-33 for more information. The `file_data` keyword matches if the HTTP Inspect preprocessor detects HTTP response body data.

- Uncompressed gzip file data

  To inspect uncompressed gzip files in the HTTP response body, the HTTP Inspect preprocessor must be enabled and you must configure the preprocessor to inspect HTTP responses and to decompress gzip-compressed files in the HTTP response body. For more information, see Decoding HTTP Traffic, page 22-31, and the **Inspect HTTP Responses** and **Inspect Compressed Data** options in Selecting Server-Level HTTP Normalization Options, page 22-33. The `file_data` keyword matches if the HTTP Inspect preprocessor detects uncompressed gzip data in the HTTP response body.

- Normalized JavaScript

  To inspect normalized JavaScript data, the HTTP Inspect preprocessor must be enabled and you must configure the preprocessor to inspect HTTP responses. See Decoding HTTP Traffic, page 22-31 and **Inspect HTTP Responses** in Selecting Server-Level HTTP Normalization Options, page 22-33 for more information. The `file_data` keyword matches if the HTTP Inspect preprocessor detects JavaScript in response body data.

- SMTP payload

  To inspect the SMTP payload, the SMTP preprocessor must be enabled. See Configuring SMTP Decoding, page 22-64 for more information. The `file_data` keyword matches if the SMTP preprocessor detects SMTP data.

- Encoded email attachments in SMTP, POP, or IMAP traffic

  To inspect email attachments in SMTP, POP, or IMAP traffic, the SMTP, POP, or IMAP preprocessor, respectively, must be enabled, alone or in any combination. Then, for each enabled preprocessor, you must ensure that the preprocessor is configured to decode each attachment encoding type that you want decoded. The attachment decoding options that you can configure for each preprocessor are: **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, and **Unix-to-Unix Decoding Depth**. See Decoding IMAP Traffic, page 22-54, Decoding POP Traffic, page 22-57, and Decoding SMTP Traffic, page 22-60 for more information.

  You can use multiple `file_data` keywords in a rule.

To point to the beginning of a specific payload type:

Step 1 On the Create Rule page, select **file_data** from the drop-down list and click **Add Option.**

The `file_data` keyword appears.

The `file_data` keyword has no arguments.

Pointing to the Beginning of the Packet Payload

License: Protection

The `pkt_data` keyword provides a pointer that serves as a reference for the positional arguments available for other keywords such as `content`, `byte_jump`, `byte_test`, and `pcre`.

When normalized FTP, telnet, or SMTP traffic is detected, the `pkt_data` keyword points to the beginning of the normalized packet payload. When other traffic is detected, the `pkt_data` keyword points to the beginning of the raw TCP or UDP payload.

The following normalization options must be enabled for the system to normalize the corresponding traffic for inspection by intrusion rules:

- To normalize FTP traffic for inspection, you must enable the FTP and Telnet preprocessor **Detect Telnet Escape codes within FTP commands** option; see Configuring Server-Level FTP Options, page 22-25.
- To normalize telnet traffic for inspection, you must enable the FTP & Telnet preprocessor **Normalize** telnet option; see Understanding Telnet Options, page 22-20.
- To normalize SMTP traffic for inspection, you must enable the SMTP preprocessor **Normalize** option; see Understanding SMTP Decoding, page 22-60.

You can use multiple `pkt_data` keywords in a rule.

To point to the beginning of the packet payload:

Step 1 On the Create Rule page, select **pkt_data** from the drop-down list and click **Add Option.**

The `pkt_data` keyword appears.

The `pkt_data` keyword has no arguments.

Decoding and Inspecting Base64 Data

License: Protection

You can use the `base64_decode` and `base64_data` keywords in combination to instruct the rules engine to decode and inspect specified data as Base64 data. This can be useful, for example, for inspecting Base64-encoded HTTP Authentication request headers and Base64-encoded data in HTTP PUT and POST requests.

These keywords are particularly useful for decoding and inspecting Base64 data in HTTP requests. However, you can also use them with any protocol such as SMTP that uses the space and tab characters the same way HTTP uses these characters to extend a lengthy header line over multiple lines. When this line extension, which is known as folding, is not present in a protocol that uses it, inspection ends at any carriage return or line feed that is not followed with a space or tab.

See the following sections for more information:

- base64_decode, page 30-100

- base64_data, page 30-100

base64_decode

License: Protection

The `base64_decode` keyword instructs the rules engine to decode packet data as Base64 data. Optional arguments let you specify the number of bytes to decode and where in the data to begin decoding.

You can use the `base64_decode` keyword once in a rule; it must precede at least one instance of the `base64_data` keyword. See base64_data, page 30-100 for more information.

Before decoding Base64 data, the rules engine unfolds lengthy headers that are folded across multiple lines. Decoding ends when the rules engine encounters any the following:

- the end of a header line
- the specified number of bytes to decode
- the end of the packet

The following table describes the arguments you can use with the `base64_decode` keyword.

Table 30-60 Optional base64_decode Arguments

Argument	Description
Bytes	Specifies the number of bytes to decode. When not specified, decoding continues to the end of a header line or the end of the packet payload, whichever comes first. You can specify a positive, non-zero value.
Offset	Determines the offset relative to the start of the packet payload or, when you also specify **Relative**, relative to the current inspection location. You can specify a positive, non-zero value.
Relative	Specifies inspection relative to the current inspection location.

To decode Base64 data:

Step 1 On the Create Rule page, select **base64_decode** from the drop-down list and click **Add Option.**

The `base64_decode` keyword appears.

Step 2 Optionally, select any of the arguments described in the Optional base64_decode Arguments table.

base64_data

License: Protection

The `base64_data` keyword provides a reference for inspecting Base64 data decoded using the `base64_decode` keyword. The `base64_data` keyword sets inspection to begin at the start of the decoded Base64 data. Optionally, you can then use the positional arguments available for other keywords such as `content` or `byte_test` to further specify the location to inspect.

You must use the `base64_data` keyword at least once after using the `base64_decode` keyword; optionally, you can use `base64_data` multiple times to return to the beginning of the decoded Base64 data.

Note the following when inspecting Base64 data:

- You cannot use the fast pattern matcher; see Use Fast Pattern Matcher, page 30-26 for more information.
- If you interrupt Base64 inspection in a rule with an intervening HTTP content argument, you must insert another `base64_data` keyword in the rule before further inspecting Base64 data; see HTTP Content Options, page 30-23 for more information.

To inspect decoded Base64 data:

Step 1 On the Create Rule page, select **base64_data** from the drop-down list and click **Add Option.**

The `base64_data` keyword appears.

Constructing a Rule

License: Protection

Just as you can create your own custom standard text rules, you can also modify existing standard text rules and shared object rule provided by Cisco and save your changes as a new rule. Note that for shared object rules provided by Cisco, you are limited to modifying rule header information such as the source and destination ports and IP addresses. You cannot modify the rule keywords and arguments in a shared object rule.

See the following sections for more information:

- Writing New Rules, page 30-101
- Modifying Existing Rules, page 30-103
- Adding Comments to Rules, page 30-104
- Deleting Custom Rules, page 30-105

Writing New Rules

License: Protection

You can create your own standard text rules.

In a custom standard text rule, you set the rule header settings and the rule keywords and arguments. Optionally, you can use the rule header settings to focus the rule to only match traffic using a specific protocol and traveling to or from specific IP addresses or ports.

After you create a new rule, you can find it again quickly using the rule number, which has the format `GID:SID:Rev`. The rule number for all standard text rules starts with 1. The second part of the rule number, the Snort ID (SID) number, indicates whether the rule is a local rule or a rule provided by Cisco. When you create a new rule, the system assigns the rule the next available Snort ID number for a local rule and saves the rule in the local rule category. Snort ID numbers for local rules start at 1,000,000 and the SID for each new local rule is incremented by one. The last part of the rule number is the revision number. For a new rule, the revision number is one. Each time you modify a custom rule the revision number increments by one.

Note The system assigns a new SID to any custom rule in an intrusion policy that you import. For more information, see Importing and Exporting Configurations, page B-1.

To write a custom standard text rule using the rule editor:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy > Rule Editor**.

The Rule Editor page appears.

Step 2 Click **Create Rule**.

The Create Rule page appears.

Step 3 In the **Message** field, enter the message you want displayed with the event.

For details on event messages, see Defining the Event Message, page 30-11.

Tip You must specify a rule message. Also, the message cannot consist of white space only, one or more quotation marks only, one or more apostrophes only, or any combination of just white space, quotation marks, or apostrophes.

Step 4 From the **Classification** list, select a classification to describe the type of event.

For details on available classifications, see Defining the Intrusion Event Classification, page 30-12.

Step 5 From the **Action** list, select the type of rule you would like to create. You can use one of the following:

- Select **alert** to create a rule that generates an event when traffic triggers the rule.
- Select **pass** to create a rule that ignores traffic that triggers the rule.

Step 6 From the **Protocol** list, select the traffic protocol (**tcp**, **udp**, **icmp**, or **ip**) of packets you want the rule to inspect.

For more information about selecting a protocol type, see Specifying Protocols, page 30-4.

Step 7 In the **Source IPs** field, enter the originating IP address or address block for traffic that should trigger the rule. In the **Destination IPs** field, enter the destination IP address or address block for traffic that should trigger the rule.

For more detailed information about the IP address syntax that the rule editor accepts, see Specifying IP Addresses In Intrusion Rules, page 30-5.

Step 8 In the **Source Port** field, enter the originating port numbers for traffic that should trigger the rule. In the **Destination Port** field, enter the receiving port numbers for traffic that should trigger the rule.

Note The system ignores port definitions in an intrusion rule header when the protocol is set to `ip`.

For more detailed information about the port syntax that the rule editor accepts, see Defining Ports in Intrusion Rules, page 30-8.

Step 9 From the **Direction** list, select the operator that indicates which direction of traffic you want to trigger the rule. You can use one of the following:

- **Directional** to match traffic that moves from the source IP address to the destination IP address
- **Bidirectional** to match traffic that moves in either direction

Step 10 From the **Detection Options** list, select the keyword that you want to use.

Step 11 Click **Add Option**.

Step 12 Enter any arguments that you want to specify for the keyword you added. For more information about rule keywords and how to use them, see Understanding Keywords and Arguments in Rules, page 30-9.

When adding keywords and arguments, you can also perform the following:

- To reorder keywords after you add them, click the up or down arrow next to the keyword you want to move.
- To delete a keyword, click the **X** next to that keyword.

Repeat steps 10 through 12 for each keyword option you want to add.

Step 13 Click **Save As New** to save the rule.

The system assigns the rule the next available Snort ID (SID) number in the rule number sequence for local rules and saves it in the local rule category.

The system does not begin evaluating traffic against new or changed rules until you enable them within the appropriate intrusion policy, and then apply the intrusion policy as part of an access control policy. See Deploying Configuration Changes, page 4-11 for more information.

Modifying Existing Rules

License: Protection

You can modify custom standard text rules. You can also modify a standard text rule or shared object rule provided by Cisco and create one or more new instances of the rule by saving it.

Creating a rule or modifying a Cisco rule copies the new rule or revision to the local rule category and assigns the rule the next available Snort ID (SID) greater than 100000.

You can only modify header information for a shared object rule. You cannot modify the rule keywords used in a shared object rule or their arguments. Modifying header information for a shared object rule and saving your changes creates a new instance of the rule with a generator ID (GID) of 3 and the next available SID for a custom rule. The Rule Editor links the new instance of the shared object rule to the reserved `soid` keyword, which maps the rule you create to the rule created by the VRT. You can delete instances of a shared object rule that you create, but you cannot delete shared object rules provided by Cisco. See Understanding Rule Headers, page 30-3 and Deleting Custom Rules, page 30-105 for more information.

Note Do not modify the protocol for a shared object rule; doing so would render the rule ineffective.

To modify a rule:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy > Rule Editor**.

The Rule Editor page appears.

Step 2 Locate the rule or rules you want to modify. You have the following options:

- To locate rules by browsing rule categories, navigate through the folders to the rule you want and click the edit icon ( ) next to the rule.

- To locate a rule or rules by filtering the rules displayed on the page, enter a rule filter in the text box indicated by the filter icon ( ) at the upper left of the rule list. Navigate to the rule you want and click the edit icon ( ) next to the rule. See Filtering Rules on the Rule Editor Page, page 30-105 for more information.

The rule editor opens, displaying the rule you selected.

Note that if you select a shared object rule, the rule editor displays only the rule header information. A shared object rule can be identified on the Rule Editor page by a listing that begins with the number 3 (the GID), for example, 3:1000004.

Step 3 Make any modifications to the rule (see Writing New Rules, page 30-101 for more information about rule options) and click **Save As New**.

The rule is saved to the local rule category.

Tip If you want to use the local modification of the rule instead of the system rule, deactivate the system rule by using the procedures at Setting Rule States, page 27-19 and activate the local rule.

Step 4 Activate the intrusion policy by applying it as part of an access control policy as described in Deploying Configuration Changes, page 4-11 to apply your changes.

Adding Comments to Rules

License: Protection

You can add comments to any intrusion rule. This allows you to provide additional context and information about the rule and the exploit or policy violation it identifies.

To add a comment to a rule:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy > Rule Editor**.

The Rule Editor page appears.

Step 2 Locate the rule you want to annotate. You have the following options:

- To locate a rule by browsing rule categories, navigate through the folders to the rule you want and click the edit icon ( ) next to the rule.
- To locate a rule by filtering the rules displayed on the page, enter a rule filter in the text box, which is indicated by the filter icon ( ), at the upper left of the rule list. Navigate to the rule you want and click the edit icon ( ) next to the rule. See Filtering Rules on the Rule Editor Page, page 30-105 for more information.

The rule editor appears.

Step 3 Click **Rule Comment**.

The Rule Comment page appears.

Step 4 Enter your comment in the text box and click **Add Comment**.

The comment is saved in the comment text box.

Deleting Custom Rules

License: Protection

You can delete custom rules that are not currently enabled in an intrusion policy. You cannot delete either standard text rules or shared object rules rules provided by Cisco.

The system stores deleted rules in the deleted category, and you can use a deleted rule as the basis for a new rule. See Modifying Existing Rules, page 30-103 for information on editing rules.

The Rules page in an intrusion policy does not display the deleted category, so you cannot enable deleted custom rules.

Note that you can also delete all local rules on the Rule Updates page. See, for example, Using One-Time Rule Updates, page 46-10.

See the following sections for more information:

- For information on creating custom rules, see Writing New Rules, page 30-101.
- For information on importing local rules, see Importing Rule Updates and Local Rule Files, page 46-9.
- For information on setting rule states, see Setting Rule States, page 27-19.

To delete custom rules:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy > Rule Editor**.

The Rule Editor page appears.

Step 2 You have two choices:

- Click **Delete Local Rules**, then click **OK**.

  All rules not currently enabled in an intrusion policy whose changes you have saved are deleted from the local rule category and moved to the deleted category.

- Navigate through the folders to the local rule category; click on the local rule category to expand it, then click the delete icon ( ) next to a rule you want to delete.

  The rule is deleted from the local rule category and moved to the deleted category.

  Note that custom standard text rules have a generator ID (GID) of 1 (for example, 1:1000012) and custom shared object rules have a GID of 3 (for example, 3:1000005).

Tip The system also stores shared object rules that you save with modified header information in the local rule category and lists them with a GID of 3. You can delete your modified version of a shared object rule, but you cannot delete the original shared object rule.

Filtering Rules on the Rule Editor Page

License: Protection

You can filter the rules on the Rule Editor page to display a subset of rules. This can be useful, for example, when you want to modify a rule or change its state but have difficulty finding it among the thousands of rules available.

When you enter a filter, the page displays any folder that includes at least one matching rule, or a message when no rule matches. Your filter can include special keywords and their arguments, character strings, and literal character strings in quotes, with spaces separating multiple filter conditions. A filter cannot include regular expressions, wild card characters, or any special operator such as a negation character (!), a greater than symbol (>), less than symbol (<), and so on.

All keywords, keyword arguments, and character strings are case-insensitive. Except for the `gid` and `sid` keywords, all arguments and strings are treated as partial strings. Arguments for `gid` and `sid` return only exact matches.

Optionally, you can expand a folder on the original, unfiltered page and the folder remains expanded when the subsequent filter returns matches in that folder. This can be useful when the rule you want to find is in a folder that contains a large number of rules.

You cannot constrain a filter with a subsequent filter. Any filter you enter searches the entire rules database and returns all matching rules. When you enter a filter while the page still displays the result of a previous filter, the page clears and returns the result of the new filter instead.

You can use the same features with rules in a filtered or unfiltered list. For example, you can edit rules in a filtered or unfiltered list on the Rule Editor page.

See the following sections for more information:

- Using Keywords in a Rule Filter, page 30-106
- Using Character Strings in a Rule Filter, page 30-107
- Combining Keywords and Character Strings in a Rule Filter, page 30-108
- Filtering Rules, page 30-108

Using Keywords in a Rule Filter

License: Protection

Each rule filter can include one or more keywords in the format:

`keyword:argument`

where *`keyword`* is one of the keywords in the Rule Filter Keywords table and *`argument`* is a single, case-insensitive, alphanumeric string to search for in the specific field or fields relevant to the keyword.

Arguments for all keywords except `gid` and `sid` are treated as partial strings. For example, the argument `123` returns `"12345"`, `"41235"`, `"45123"`, and so on. The arguments for `gid` and `sid` return only exact matches; for example, `sid:3080` returns only SID 3080.

Tip You can search for a partial SID by filtering with one or more character strings. See Using Character Strings in a Rule Filter, page 30-107 for more information.

The following table describes the specific filtering keywords and arguments you can use to filter rules.

Table 30-61 Rule Filter Keywords

Keyword	Description	Example
`arachnids`	Returns one or more rules based on all or part of the Arachnids ID in a rule reference. See Defining the Event Reference, page 30-14 for more information.	`arachnids:181`
`bugtraq`	Returns one or more rules based on all or part of the Bugtraq ID in a rule reference. See Defining the Event Reference, page 30-14 for more information.	`bugtraq:2120`
`cve`	Returns one or more rules based on all or part of the CVE number in a rule reference. See Defining the Event Reference, page 30-14 for more information.	`cve:2003-0109`
`gid`	The argument `1` returns standard text rules. The argument `3` returns shared object rules. See Table 27-1 on page 27-2 for more information.	`gid:3`
`mcafee`	Returns one or more rules based on all or part of the McAfee ID in a rule reference. See Defining the Event Reference, page 30-14 for more information.	`mcafee:10566`
`msg`	Returns one or more rules based on all or part of the rule Message field, also known as the event message. See Defining the Event Message, page 30-11 for more information.	`msg:chat`
`nessus`	Returns one or more rules based on all or part of the Nessus ID in a rule reference. See Defining the Event Reference, page 30-14 for more information.	`nessus:10737`
`ref`	Returns one or more rules based on all or part of a single alphanumeric string in a rule reference or in the rule Message field. See Defining the Event Reference, page 30-14 and Defining the Event Message, page 30-11 for more information.	`ref:MS03-039`
`sid`	Returns the rule with the exact Signature ID.	`sid:235`
`url`	Returns one or more rules based on all or part of the URL in a rule reference. See Defining the Event Reference, page 30-14 for more information.	`url:faqs.org`

Using Character Strings in a Rule Filter

License: Protection

Each rule filter can include one or more alphanumeric character strings. Character strings search the rule **Message** field, Signature ID, and Generator ID. For example, the string `123` returns the strings `"Lotus123"`, `"123mania"`, and so on in the rule message, and also returns SID 6123, SID 12375, and so on. For information on the rule **Message** field, see Defining the Event Message, page 30-11.

All character strings are case-insensitive and are treated as partial strings. For example, any of the strings `ADMIN`, `admin`, or `Admin` return `"admin"`, `"CFADMIN"`, `"Administrator"` and so on.

You can enclose character strings in quotes to return exact matches. For example, the literal string `"overflow attempt"` in quotes returns only that exact string, whereas a filter comprised of the two strings `overflow` and `attempt` without quotes returns `"overflow attempt"`, `"overflow multipacket attempt"`, `"overflow with evasion attempt"`, and so on.

Combining Keywords and Character Strings in a Rule Filter

License: Protection

You can narrow filter results by entering any combination of keywords, character strings, or both, separated by spaces. The result includes any rule that matches all the filter conditions.

You can enter multiple filter conditions in any order. For example, each of the following filters returns the same rules:

- `url:at login attempt cve:200`
- `login attempt cve:200 url:at`
- `login cve:200 attempt url:at`

Filtering Rules

License: Protection

You can filter the rules on the Rule Editor page to display a subset of rules so you can more easily find specific rules. You can then use any of the page features.

To filter for specific rules:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy > Rule Editor**.

The Rule Editor page appears.

Rule filtering can be particularly useful on the Rule Editor page when you want to locate a rule to edit it. See Modifying Existing Rules, page 30-103 for more information.

Step 2 Optionally, select a different grouping method from the Group Rules By list.

Tip Filtering may take significantly longer when the combined total of rules in all sub-groups is large because rules appear in multiple categories, even when the total number of unique rules is much smaller.

Step 3 Optionally, click the folder next to any group that you want to expand.

The folder expands to show the rules in that group. Note that some rule groups have sub-groups that you can also expand.

Note also that expanding a group on the original, unfiltered page can be useful when you expect that a rule might be in that group. The group remains expanded when the subsequent filter results in a match in that folder, and when you return to the original, unfiltered page by clicking on the filter clearing icon (✖).

Step 4 To activate the filter text box, click to the right of the filter icon ( ) that is inside the text box at the upper left of the rule list.

Step 5 Type your filter constraints and press Enter.

Your filter can include keywords and arguments, character strings with or without quotes, and spaces separating multiple conditions. See Filtering Rules on the Rule Editor Page, page 30-105 for more information.

The page refreshes to display any group that contains at least one matching rule.

Step 6 Optionally, open any folder not already opened to display matching rules. You have the following filtering choices:

- To enter a new filter, position your cursor inside the filter text box and click to activate it; type your filter and press Enter.
- To clear the current filtered list and return to the original, unfiltered page, click the filter clearing icon (✖).

Step 7 Optionally, make any changes to the rule that you would normally make on the page. See Modifying Existing Rules, page 30-103.

To put any changes you make into effect, apply the intrusion policy part of an access control policy as described in Deploying Configuration Changes, page 4-11.



Introduction to Identity Data

You can configure identity policies to use User Agents, ISE/ISE-PIC devices, or captive portal to obtain data about the users on your network. For more information, see User Identity Sources, page 33-1.

Uses for Identity Data

Collecting identity data allows you to take advantage of many features, including:

- perform user control by writing access control rules using realm, user, user group, and ISE attribute conditions
- alert you via email, SNMP trap, or syslog when the system generates an intrusion event with a specific impact flag

User Detection Fundamentals

You can use your identity policies to monitor user activity on your network, which allows you to correlate threat, endpoint, and network intelligence with user identity information. By linking network behavior, traffic, and events directly to individual users, the system can help you to identify the source of policy breaches, attacks, or network vulnerabilities. For example, you could determine:

- who owns the host targeted by an intrusion event that has a Vulnerable (level 1: red) impact level
- who initiated an internal attack or portscan
- who is attempting unauthorized access of a server that has high host criticality
- who is consuming an unreasonable amount of bandwidth
- who has not applied critical operating system updates
- who is using instant messaging software or peer-to-peer file-sharing applications in violation of company IT policy

Armed with this information, you can use other features of the ASA FirePOWER module to mitigate risk, perform access control, and take action to protect others from disruption. These capabilities also significantly improve audit controls and enhance regulatory compliance.

After you configure user identity sources, you can perform user awareness and user control.

User awareness

The ability to view and analyze user data

User control

The ability to configure user access control rule conditions to block users or user activity in traffic on your network, based on conclusions you drew from user awareness.

You can obtain user data from authoritative identity sources (referenced by your identity policy).

An identity source is authoritative if a trusted server validated the user login. You can use the data obtained from authoritative logins to perform user awareness and user control. Authoritative user logins are obtained from passive and active authentications:

- *Passive authentications* occur when a user authenticates through an external server. The User Agent and ISE/ISE-PIC are the only passive authentication methods supported by the ASA FirePOWER module.
- *Active authentications* occur when a user authenticates through a Firepower device. Captive portal is the only active authentication method supported by the ASA FirePOWER module.

The following table provides a brief overview of the user identity sources supported by the ASA FirePOWER module.

Table 31-1

User Identity Source	Server Requirements	Source Type	Authentication Type	User Awareness?	User Access Control?	For more information, see...
User Agent	Microsoft Active Directory	authoritative logins	passive	Yes	Yes	The User Agent Identity Source, page 33-3
ISE/ISE-PIC	Microsoft Active Directory	authoritative logins	passive	Yes	Yes	The ISE/ISE-PIC Identity Source, page 33-4
Captive portal	LDAP or Microsoft Active Directory	authoritative logins	active	Yes	Yes	The Captive Portal Active Authentication Identity Source, page 33-7

Consider the following when selecting identity sources to deploy:

- you must use captive portal to record failed authentication activity. A failed authentication attempt does not add a new user to the list of users in the database.
- you must deploy an appliance that has an IP address for its sensing interface (for example, a routed interface) in order to use captive portal.

User Identity Deployments

When the system detects user data from a user login, from any identity source, the user from the login is checked against the list of users in the user database. If the login user matches an existing user, the data from the login is assigned to the user. Logins that do not match existing users cause a new user to be created, unless the login is in SMTP traffic. Non-matching logins in SMTP traffic are discarded.

The User Activity Database

The user activity database on the device contains records of user activity on your network reported by all of your configured identity sources. The system logs events in the following circumstances:

- when it detects individual logins or logoffs
- when it detects a new user
- when you manually delete a user
- when the system detects a user that is not in the database, but cannot add the user because you have reached your user limit

The Users Database

The users database contains a record for each user reported by your configured identity sources.

The total number of users the device can store depends on the model. When the limit has been reached, you must delete users (manually or with a database purge) to allow new users to be added.

If an identity source is configured to exclude specific user names, user activity data for those user names are not reported to the ASA FirePOWER module. These excluded user names remain in the database, but are not associated with IP addresses.

Current User Identities

When the system detects multiple logins to the same host by different users, the system assumes that only one user is logged into any given host at a time, and that the current user of a host is the last authoritative user login. If multiple users are logged in through remote sessions, the last user reported by the server is the user reported to the ASA FirePOWER module.

When the system detects multiple logins to the same host by the same user, the system records the first time that a user logs into a specific host and disregards subsequent logins. If an individual user is the only person who logs into a specific host, the only login that the system records is the original login.

If another user logs into that host, however, the system records the new login. Then, if the original user logs in again, his or her new login is recorded.

User Database Limits

Your device model determines how many users you can monitor, as well as how many users you can use to perform user control.

When deploying an ASA FirePOWER module managed via ASDM, you can store a maximum of 2,000 authoritative users in the Users database.



Realms and Identity Policies

A *realm* consists of one or more LDAP or Microsoft Active Directory servers that share the same credentials. You must configure a realm if you want to perform user and user group queries, user access control, or to configure a User Agent, ISE/ISE-PIC, or captive portal. After configuring one or more realms, you can configure an identity policy.

An *identity policy* associates traffic on your network with an authoritative identity source and a realm. After configuring your identity policy, you can associate it with an access control policy and deploy the access control policy to your device.

Realm Fundamentals

License: Any

Realms establish connections between the ASA FirePOWER module and the servers targeted for monitoring. They specify the connection settings and authentication filter settings for the server. Realms can:

- specify the users and user groups whose activity you want to monitor.
- allow you to query the server for user metadata on authoritative users.

You can add multiple servers as directories within a realm, but they must share the same basic realm information. The directories within a realm must be exclusively LDAP or exclusively AD servers. After you enable a realm, your saved changes take effect next time the ASA FirePOWER module queries the server.

To perform user awareness, you must configure a realm for any of the supported server types. The module uses these connections to query the servers for data associated with POP3 and IMAP users. The module uses the email addresses in POP3 and IMAP logins to correlate with LDAP users on an Active Directory, OpenLDAP, or Oracle Directory Server Enterprise Edition server. For example, if a device detects a POP3 login for a user with the same email address as an LDAP user, the module associates the LDAP user's metadata with that user.

To perform user access control, you can configure the following:

- a realm for an AD server configured for either a User Agent or ISE/ISE-PIC device.

Note Configuring a realm is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions.

- a realm for an Oracle or OpenLDAP server configured for captive portal.

If you configure a realm to download users (for user awareness or user control), the ASA FirePOWER module regularly queries the server to obtain metadata for new and updated users whose activity was detected since the last query.

User activity data is stored in the user activity database and user identity data is stored in the users database. The maximum number of users you can store and use in access control depends on your device model. When choosing which users and groups to include, make sure the total number of users is less than your model limit. If your access control parameters are too broad, the ASA FirePOWER module obtains information on as many users as it can and reports the number of users it failed to retrieve in the task queue.



Note If you remove a user that has been detected by the module from your LDAP servers, the ASA FirePOWER module does not remove that user from its users database; you must manually delete it. However, your LDAP changes are reflected in access control rules when the ASA FirePOWER module next updates its list of authoritative users.

Supported Servers for Realms

License: Any

You can configure realms to connect to the following types of servers, providing they have TCP/IP access from the ASA FirePOWER module:

Table 32-1 Supported Servers for Realms

Server Type	Supported for user awareness data retrieval?	Supported for User Agent data retrieval?	Supported for ISE/ISE-PIC data retrieval?	Supported for captive portal data retrieval?
Microsoft Active Directory on Windows Server 2003, Windows Server 2008, and Windows Server 2012	Yes	Yes	Yes	Yes, except Windows Server 2003 if you are using NTLM captive portal
Oracle Directory Server Enterprise Edition 7.0 on Windows Server 2003 and Windows Server 2008	Yes	No	No	Yes
OpenLDAP on Linux	Yes	No	No	Yes

Note the following about your server group configurations:

- If you want to perform user control on user groups or on users within groups, you must configure user groups on the server. The ASA FirePOWER module cannot perform user group control if the server organizes the users in basic object hierarchy.

  Cisco recommends that you limit the size of your LDAP or AD server groups to contain a maximum of 1500 users. Configuring realms to include or exclude oversized groups, or creating access control rules that target oversized user groups may result in performance issues.
- By default, AD servers limit the number of users they report from secondary groups. You must customize this limit so that all of the users in your secondary groups are reported to the ASA FirePOWER module.

Supported Server Field Names

License: Any

The servers in your realms must use the field names listed in the following table in order for the ASA FirePOWER module to retrieve user metadata the servers. If the field names are incorrect on your server, the ASA FirePOWER module cannot populate its database with the information in that field.

Table 32-2 Mapping Server Fields to ASA FirePOWER Fields

Metadata	ASA FirePOWER module	Active Directory	Oracle Directory Server	OpenLDAP
LDAP user name	Username	samaccountname	cn uid	cn uid
first name	First Name	givenname	givenname	givenname
last name	Last Name	sn	sn	sn
email address	Email	mail userprincipalname (if mail has no value)	mail	mail
department	Department	department distinguishedname (if department has novalue)	department	ou
telephone number	Phone	telephonenumber	n/a	telephonenumber

Troubleshooting Issues with Realms

License: Any

If you notice unexpected server connection behavior, consider tuning your realm configuration, device settings, or server settings.

User timeouts are occurring at unexpected times

If you notice the system performing user timeouts at unexpected intervals, confirm that the time on your User Agent or ISE/ISE-PIC device is synchronized with the time on the ASA FirePOWER module. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.

Users are not included or excluded as specified in your realm configuration

If you configure a realm for an Active Directory server that includes or excludes users who are members of a secondary group on your Active Directory server, your server may be limiting the number of users it reports.

By default, Active Directory servers limit the number of users they report from secondary groups. You must customize this limit so that all of the users in your secondary groups are reported to the ASA FirePOWER module.

User download is slow

If you notice that user download is slow, confirm that your LDAP and AD server groups contain a maximum of 1500 users. Configuring realms to include or exclude oversized user groups may result in performance issues.

Identity Policy Fundamentals

License: Any

Identity policies contain identity rules. Identity rules associate sets of traffic with a realm and an authentication method: passive authentication, active authentication, or no authentication.

You must fully configure the realms and authentication methods you plan to use before you can invoke them in your identity rules:

- You configure realms outside of your identity policy, at **Configuration > ASA FirePOWER Configuration > Integration > Realms**.
- You configure the passive authentication identity sources, the User Agent and ISE/ISE-PIC, at **Configuration > ASA FirePOWER Configuration > Integration > Identity Sources**.
- You configure the active authentication identity source, captive portal, within the identity policy.

After you configure one or more identity policies, you must invoke one identity policy in your access control policy. When traffic on your network matches the conditions in your identity rule and the authentication method is passive or active, the module associates the traffic with the specified realm and authenticates the users in the traffic using the specified identity source.

If you do not configure an identity policy, the module does not perform user authentication.

Creating a Realm

License: Control

To create a realm:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Integration**.

Step 2 Click **Realms**.

Step 3 Click **New Realm**.

Step 4 Configure basic realm information as described in Configuring Basic Realm Information, page 32-7.

Step 5 Configure directories as described in Configuring a Realm Directory, page 32-7.

Step 6 Configure user and user group download (required for access control) as described in Configuring Automatic User Download, page 32-8.

Step 7 Save the realm settings.

Step 8 Optionally, edit the realm and modify the default User Session Timeout settings as described in Configuring Realm User Session Timeouts, page 32-8.

Step 9 Save the realm settings.

What to Do Next

- Enable the realm as described in Enabling or Disabling a Realm, page 32-18.
- Optionally, monitor the task status; see the Task Status page (**Monitoring > ASA FirePOWER Monitoring > Task Status**).

Realm Fields

License: Any

The following fields are used to configure a realm.

Realm Configuration Fields

AD Primary Domain

For AD realms only, the domain for the Active Directory server where users should be authenticated.

AD Join Username and AD Join Password

For AD realms intended for Kerberos captive portal active authentication, the distinguished username and password for a user with appropriate rights to join clients to the domain.

If you select Kerberos (or HTTP Negotiate, if you want Kerberos as an option) as the **Authentication Type** in an identity rule, the **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** in order to perform Kerberos captive portal authentication.

Description

An optional description for the realm.

Directory Username and Directory Password

The distinguished username and password for a user with appropriate rights to the user information you want to retrieve.

Base DN

The directory tree on the server where the ASA FirePOWER module should begin searching for user data.

Typically, the base DN has a basic structure indicating the company domain and operational unit. For example, the Security organization of the Example company might have a base DN of ou=security,dc=example,dc=com.

Group DN

The directory tree on the server where the ASA FirePOWER module should search for users with the group attribute.

Group Attribute

The group attribute for the server: Member, Unique Member, or Custom.

Name

A unique name for the realm.

Type

The type of realm, AD or LDAP.

User Session Timeout: Authenticated Users

The maximum amount of time, in minutes, before a user's session is timed out.

If a user was passively authenticated and their session times out, they are identified as Unknown and their current session is allowed or blocked depending on their access control rule settings. The module re-identifies the user the next time they log in.

If a user was actively authenticated (captive portal) and their session times out, they are prompted to re-authenticate.

User Session Timeout: Failed Authentication Users

The amount of time, in minutes, after a failed active authentication attempt that a user's session is timed out. When a user fails to authenticate and their session times out, they are prompted to re-authenticate.

User Session Timeout: Guest Users

The maximum amount of time, in minutes, before an actively authenticated (captive portal) guest user's session is timed out. When their session times out, they are prompted to re-authenticate.

Realm Directory Fields

These settings apply to individual servers (directories) within a realm.

Encryption

The encryption method you want to use for the server connection. If you specify an Encryption method, you must specify a host name in this field.

Hostname / IP Address

The hostname or IP address for the server.

Port

The port you want to use for the server connection.

SSL Certificate

The SSL certificate you want to use for authentication to the server. You must configure the **Encryption** type in order to use an SSL certificate.

If you are using a certificate to authenticate, the name of the server in the certificate must match the server **Hostname / IP Address**. For example, if you use 10.10.10.250 as the IP address but computer1.example.com in the certificate, the connection fails.

User Download Fields

Download for access control

Selecting this check box configures the automatic download of user data. You can use the data for user awareness and, in some cases, user access control.

Use the **Begin automatic download at and Repeat every drop-down** menus to configure the download frequency.

Configuring Basic Realm Information

License: Control

To configure basic realm information:

Step 1 On the Add New Realm page, type a **Name** and, optionally, a **Description**.

Step 2 Select a **Type** from the drop-down list.

Step 3 If you are configuring an AD realm, enter an **AD Primary Domain**.

Step 4 If you are configuring an AD realm intended for Kerberos captive portal active authentication, enter a distinguished **AD Join Username** and **AD Join Password** for a user with appropriate rights to join clients to the domain.

Step 5 Enter a distinguished **Directory Username** and **Directory Password** for a user with appropriate rights to the user information you want to retrieve.

Step 6 Enter a **Base DN** for the directory.

Step 7 Enter a **Group DN** for the directory.

Step 8 Optionally, select a **Group Attribute** from the drop-down list.

Step 9 Click **OK**.

What to Do Next

- Configure the realm directory as described in Configuring a Realm Directory, page 32-7.

Configuring a Realm Directory

License: Control

To configure a realm directory:

Step 1 On the Directory tab, click **Add Directory**.

Step 2 Enter the Hostname / IP Address and Port for the server.

Step 3 Select an Encryption Mode.

Step 4 Optionally, select an SSL Certificate from the drop-down list. Note that you can click the add icon ( ) to create an object on the fly.

Step 5 If you want to test the connection, click **Test**.

Step 6 Click **OK**.

What to Do Next

- Optionally, configure automatic user download as described in Configuring Automatic User Download, page 32-8.

Configuring Automatic User Download

License: Control

If you do not specify any groups to include, the ASA FirePOWER module retrieves user data for all the groups that match the parameters you provided. For performance reasons, Cisco recommends that you explicitly include only the groups that represent the users you want to use in access control.

To configure automatic user download:

Step 1 On the User Download tab, select the **Download users and groups (required for user access control)** check box.

Step 2 Select a time to **Begin automatic download at** from the drop-down lists.

Step 3 Select a download interval from the **Repeat Every** drop-down list.

Step 4 To include or exclude user groups from the download, select user groups from the **Available Groups** column and click **Add to Include** or **Add to Exclude**.

Step 5 To include or exclude individual users, type the user into the field below Groups to Include or Groups to Exclude and click **Add**.

Note Excluding users from download prevents you from writing an access control rule with that user as a condition. Separate multiple users with commas. You can also use an asterisk (*) as a wildcard character in this field.

Configuring Realm User Session Timeouts

License: Control

Note If the module is performing user timeouts at unexpected intervals, confirm that the time on your User Agent or ISE/ISE-PIC device is synchronized with the time on the ASA FirePOWER module.

To configure realm user session timeouts:

Step 1 Select the **Realm Configuration** tab.

Step 2 Enter user session timeout values for **Authenticated Users**, **Failed Authentication Users**, and **Guest Users**.

Step 3 Click **Save** or continue editing the realm.

Configuring an Identity Policy

License: Control

Before You Begin

- Create and enable one or more realms as described in Creating a Realm, page 32-4.

To configure an Identity Policy:

Access: Admin/Access Admin/Network Admin

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Identity Policy**.

Step 2 Type a **Name** and, optionally, a **Description**.

Step 3 If you want to add a rule to the policy, click Add Rule as described in Creating an Identity Rule, page 32-12

Step 4 If you want to add a rule category, click Add Category as described in Adding an Identity Rule Category, page 32-19

Step 5 If you want to configure active authentication using captive portal, click Active Authentication as described in Configuring Captive Portal (Active Authentication), page 32-10.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Captive Portal (Active Authentication) Fields

License: Any

Use the following fields to configure captive portal.

Server Certificate

The server certificate presented by the captive portal daemon.

Port

The port number you want to use for the captive portal connection. The port number in this field must match the port number you configured on the ASA FirePOWER device using the `captive-portal` CLI command.

Maximum login attempts

The maximum allowed number of failed login attempts before the module denies a user's login request.

Active Authentication Response Page

The system-provided or custom HTTP response page you want to display to captive portal users. After you select an Active Authentication Response page in your identity policy active authentication settings, you must also configure one or more identity rules with HTTP Response Page as the Authentication Type.

The system-provided HTTP response page includes Username and Password fields, as well as a **Login as guest** button to allow users to access the network as guests. If you want to display a single login method, configure a custom HTTP response page.

Configuring Captive Portal (Active Authentication)

License: Control

You can select either a system-provided or a custom HTTP response page to display to captive portal users. The system-provided HTTP response page includes Username and Password fields, as well as a **Login as guest** button to allow users to access the network as guests. If you want to display a single login method, configure a custom HTTP response page.

For more information about captive portal, see The Captive Portal Active Authentication Identity Source, page 33-7.

Before You Begin

- Confirm that your device manages one or more ASA FirePOWER devices in routed mode running Version 9.5(2) or later.
- Configure an access control rule to allow traffic destined for the port you plan to use for captive portal.
- If you want to perform active authentication via captive portal on HTTPS traffic, you must create SSL rules to decrypt the traffic originating from the users you want to authenticate using captive portal.
- If you want to decrypt traffic in the captive portal connection, create an SSL rule to decrypt the traffic destined for the port you plan to use for captive portal.
- Use the `captive-portal` ASA CLI command to enable captive portal for active authentication and define the port as described in the *ASA Firewall Configuration Guide* (Version 9.5(2) or later): http://www.cisco.com/c/en/us/support/security/asa-5500-series-next-generation-firewalls/products-installation-and-configuration-guides-list.html.

To configure captive portal:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Identity Policy** and edit an identity policy.

Step 2 Click **Active Authentication**.

Step 3 Select the appropriate **Server Certificate** from the drop-down list. Optionally, click the add icon ( ) to create an object on the fly.

Step 4 Type a **Port** and specify the **Maximum login attempts**.

Step 5 Optionally, to authenticate users through a HTTP response page, select an **Active Authentication Response Page**.

Step 6 Click **Save**.

Step 7 Configure an identity rule with **Active Authentication** as the **Action** as described in Creating an Identity Rule, page 32-12. If you selected a response page in step 5, you must also select HTTP Response Page as the **Authentication Type**.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Excluding Applications From Active Authentication

License: Control

You can select applications (identified by their HTTP User-Agent strings) and exempt them from captive portal (active authentication). This allows traffic from the selected applications to pass through the identity policy without authenticating.

To exclude applications from active authentication:

Step 1 On the **Realm & Settings** tab of the identity rule editor page, use Cisco-provided filters in the **Application Filters** list to narrow the list of applications you want to add to the filter.

- Click the arrow next to each filter type to expand and collapse the list.
- Right-click a filter type and click **Check All** or **Uncheck All**. Note that the list indicates how many filters you have selected of each type.
- To narrow the filters that appear, type a search string in the **Search by name** field; this is especially useful for categories and tags. To clear the search, click the clear icon ( ).
- To refresh the filters list and clear any selected filters, click the reload icon ( ).
- To clear all filters and search fields, click **Clear All Filters**.

Note The list displays 100 applications at a time.

Step 2 Select the applications that you want to add to the filter from the **Available Applications** list:

- Select **All apps matching the filter** to add all the applications that meet the constraints you specified in the previous step.
- To narrow the individual applications that appear, type a search string in the **Search by name** field. To clear the search, click the clear icon ( ).
- Use the paging icons at the bottom of the list to browse the list of individual available applications.
- To refresh the applications list and clear any selected applications, click the reload icon ( ).

Step 3 Add the selected applications to exclude from external authentication. You can click and drag, or you can click **Add to Rule**. The result is the combination of:

- the selected Application Filters

- either the selected individual Available Applications, or **All apps matching the filter**

What to Do Next

- Continue configuring the identity rule as described in Creating an Identity Rule, page 32-12.

Associating an Identity Policy with an Access Control Policy

License: Control

You can have one identity policy currently applied to an ASA FirePOWER module. You cannot apply an identity policy independently. You cannot delete an identity policy that has been applied or is currently applying.

To associate an Identity Policy with an Access Control Policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

Step 2 Select the **Advanced** tab.

Step 3 Click the edit icon ( ) next to Identity Policy Settings.

Step 4 Select an identity policy from the drop-down.

Step 5 Click **OK**.

Step 6 Click **Store ASA FirePOWER Changes** to save your changes.

Creating an Identity Rule

License: Control

To create an identity rule:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Identity Policy**.

Step 2 Click **Add Rule**.

Step 3 Configure basic identity rule information as described in Configuring Basic Identity Rule Information, page 32-14.

Step 4 Optionally, add a zone condition as described in Adding a Zone Condition to an Identity Rule, page 32-16.

Note If you are configuring the rule for captive portal and your captive portal device contains inline and routed interfaces, you must configure a zone condition to target only the routed interfaces on the device.

Step 5 Optionally, add a network or geolocation condition as described in Adding a Network or Geolocation Condition to an Identity Rule, page 32-15.

Step 6 Optionally, add a port condition as described in Adding a Port Condition to an Identity Rule, page 32-15.

Step 7 Associate the rule with a realm as described in Associating a Realm and Configuring Active Authentication Settings in an Identity Rule, page 32-16.

Step 8 Click **Add**.

Step 9 Click **Store ASA FirePOWER Changes**.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Identity Rule Fields

Use the following fields to configure identity rules.

Enabled

Selecting this option enables the identity rule in the identity policy. Deselecting this option disables the identity rule.

Action

The type of authentication you want to perform on the users in the specified **Realm**. You can select Passive Authentication (User Agent or ISE/ISE-PIC), Active Authentication (captive portal), or No Authentication. You must fully configure the authentication method, or identity source, before selecting it as the action in an identity rule.

Realm

The realm containing the users you want to perform the specified **Action** on. You must fully configure a realm before selecting it as the realm in an identity rule.

If you select Kerberos (or HTTP Negotiate, if you want Kerberos as an option) as the A**uthentication Type** in an identity rule, the **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** in order to perform Kerberos captive portal authentication.

Use active authentication if passive authentication cannot identify user

Selecting this option authenticates users via active authentication if passive authentication fails to identify them. You must configure active authentication (captive portal) in order to select this option.

If you disable this option, users that passive authentication cannot identify are identified as Unknown. You must set the rule action to Passive Authentication in order to see this field.

Identify as Special Identities/Guest if authentication cannot identify user

Selecting this option identifies unknown users as **Special Identities/Guest** in all areas of the ASDM interface. You must set the rule action to Active Authentication or select **Use active authentication if passive authentication cannot identify user** in order to see this field.

Authentication Type

The method you want to use to perform active authentication. The selections vary depending on the type of realm, LDAP or AD:

– Select HTTP Basic if you want to authenticate users using an unencrypted HTTP Basic Authentication (BA) connection. Users log in to the network using their browser's default authentication popup window.

Most web browsers cache the credentials from HTTP Basic logins and use the credentials to seamlessly begin a new session after an old session times out.

– Select NTLM if you want to authenticate users using a NT LAN Manager (NTLM) connection. This selection is only available when you select an AD realm. Users log in to the network using their browser's default authentication popup window. If you select NTLM as your identity rule Authentication Type, you cannot use a 2003 Windows Server as your identity rule realm.

– Select Kerberos if you want to authenticate users using a Kerberos connection. This selection is available only when you select an AD realm for a server with secure LDAP (LDAPS) enabled. If transparent authentication is configured in a user's browser, the user is automatically logged in. If transparent authentication is not configured, users log in to the network using their browser's default authentication popup window.

  The **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** in order to perform Kerberos captive portal authentication.

Note If you have DNS resolution configured and you create an identity rule to perform Kerberos (or HTTP Negotiate, if you want Kerberos as an option) captive portal, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN must match the hostname you provided when configuring DNS. For ASA with FirePOWER Services devices, the FQDN must resolve to the IP address of the routed interface used for captive portal.

– Select HTTP Negotiate to allow the captive portal server to choose between HTTP Basic, Kerberos, or NTLM for the authentication connection. This selection is only available when you select an AD realm. Users log in to the network using their browser's default authentication popup window.

  The **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** in order to perform Kerberos captive portal authentication.

  If you are creating an identity rule to perform HTTP Negotiate captive portal and you have DNS resolution configured, you must configure your DNS server to resolve the hostname of the captive portal device. The hostname of the device you are using for captive portal must match the host name you provided when configuring DNS.

– Select HTTP Response Page if you want to authenticate users using a ASA FirePOWER module-provided or custom HTTP response page. Users log in to the network using the response page you configure.

  The system-provided HTTP response page includes Username and Password fields, as well as a **Login as guest** button to allow users to access the network as guests. If you want to display a single login method, configure a custom HTTP response page.

  Users who log in as guests appear in the web interface with the username **Guest**, and their realm is the realm specified in the identity rule.

Configuring Basic Identity Rule Information

License: Control

To configure basic identity rule information:

Step 1 On the identity rule editor page, type a **Name**.

Step 2 Specify whether the rule is **Enabled**.

Step 3 To add the rule to a rule category, see Adding an Identity Rule Category, page 32-19.

Step 4 Select a rule **Action** from the drop-down list.

Step 5 Click **Add** or continue editing the rule.

Adding a Network or Geolocation Condition to an Identity Rule

License: Control

To add a network or geolocation condition to an Identity Rule:

Step 1 On the identity rule editor page, select the **Networks** tab.

Step 2 Find the networks you want to add from the **Available Networks**, as follows:

- To add a network object on the fly, which you can then add to the condition, click the add icon ( ) above the Available Networks list.
- To search for network or geolocation objects to add, select the appropriate tab, click the **Search by name or value** prompt above the **Available Networks** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.

Step 3 To select an object, click it. To select all objects, right-click and then select **Select All**.

Step 4 Click **Add to Source** or **Add to Destination**.

Step 5 Add any source or destination IP addresses or address blocks that you want to specify manually. Click the **Enter an IP address** prompt below the **Source Networks** or **Destination Networks** list; then type an IP address or address block and click **Add**.

Step 6 Click **Add** or continue editing the rule.

Adding a Port Condition to an Identity Rule

License: Control

To add a port condition to an Identity Rule:

Step 1 On the identity rule editor page, select the **Ports** tab.

Step 2 Find the TCP ports you want to add from the **Available Ports**, as follows:

- To add a TCP port object on the fly, which you can then add to the condition, click the add icon ( ) above the Available Ports list.
- To search for TCP-based port objects and groups to add, click the **Search by name or value** prompt above the **Available Ports** list, then type either the name of the object, or the value of a port in the object. The list updates as you type to display matching objects. For example, if you type 443, the ASA FirePOWER module displays the provided HTTPS port object.

Step 3 To select a TCP-based port object, click it. To select all TCP-based port objects, right-click and then select **Select All**. If the object includes non-TCP-based ports, you cannot add it to your port condition.

Step 4 Click **Add to Source** or **Add to Destination**.

Step 5 Enter a **Port** under the **Selected Source Ports** or **Selected Destination Ports** list to manually specify source or destination ports. You can specify a single port with a value from 0 to 65535.

Step 6 Click **Add**.

Note The ASA FirePOWER module will not add a port to a rule condition that results in an invalid configuration.

Step 7 Click **Add** or continue editing the rule.

Adding a Zone Condition to an Identity Rule

License: Control

If the device you want to use for captive portal contains both inline and routed interfaces, you must configure a zone condition in your captive portal identity rules to target only the routed interfaces on the captive portal device.

For more information about security zones, see Working with Security Zones, page 2-32.

To add a Zone Condition to an Identity Rule:

Step 1 On the identity rule editor page, select the **Zones** tab.

Step 2 Find the zones you want to add from the **Available Zones**. To search for zones to add, click the **Search by name** prompt above the **Available Zones** list, then type a zone name. The list updates as you type to display matching zones.

Step 3 Click to select a zone. To select all zones, right-click and then select **Select All**.

Step 4 Click **Add to Source** or **Add to Destination**.

Step 5 Click **Add** or continue editing the rule.

Associating a Realm and Configuring Active Authentication Settings in an Identity Rule

License: Control

Associate the identity rule with a realm and, optionally, configure additional settings for active authentication.

To associate Identity Rules With a Realm:

Step 1 On the identity rule editor page, select the **Realm & Settings** tab.

Step 2 Select a **Realm** from the drop-down list.

Step 3 Optionally, select the **Use active authentication if passive authentication cannot identify user** check box. Note that this check box appears only when configuring a Passive Authentication rule.

Step 4 If you selected the check box in step 3, or if this is an Active Authentication rule, continue with step 4. Otherwise, skip to step 8.

Step 5 Optionally, select the **Identify as Special Identities/Guest if authentication cannot identify user** check box.

Step 6 Select an **Authentication Type** from the drop-down list.

Step 7 Optionally, **Exclude HTTP User-Agents** to exempt specific application traffic from active authentication as described in Excluding Applications From Active Authentication, page 32-11.

Step 8 Click **Add** or continue editing the rule.

Managing Realms

License: Control

To manage a Realm:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Integration > Realms**.

Step 2 If you want to delete a realm, click the delete icon ( ).

Step 3 If you want to edit a realm, click the edit icon ( ) next to the realm and make changes as described in Creating a Realm, page 32-4.

Step 4 If you want to enable or disable a realm, click the State slider next to the realm you want to enable or disable as described in Enabling or Disabling a Realm, page 32-18.

Step 5 If you want to download users and user groups on demand, click the download icon ( ) as described in Downloading Users and User Groups On-Demand, page 32-18.

Step 6 If you want to copy a realm, click the copy icon ( ).

Step 7 If you want to compare realms, see Comparing Realms, page 32-17.

Comparing Realms

License: Control

To Compare Realms:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Integration > Realms**.

Step 2 Click **Compare Realms**.

Step 3 Select **Compare Realm** from the **Compare Against** drop-down list.

Step 4 Select the realms you want to compare from the **Realm A** and **Realm B** drop-down lists.

Step 5 Click **OK**.

Step 6 If you want to navigate individually through changes, click **Previous** or **Next** above the title bar.

Step 7 Optionally, click **Comparison Report** to generate the realm comparison report.

Step 8 Optionally, click **New Comparison** to generate a new realm comparison view.

Downloading Users and User Groups On-Demand

License: Control

If you change the user or group download parameters in a realm, or if you change the users or groups on your server and want the changes to be immediately available for user control, you can force the ASA FirePOWER module to perform an on-demand user download from the server.

The maximum number of users the ASA FirePOWER module can retrieve from the server depends on your device model. If the download parameters in your realm are too broad, the ASA FirePOWER module obtains information on as many users as it can and reports the number of users it failed to retrieve in the task queue.

Before You Begin

- Enable the realm as described in Enabling or Disabling a Realm, page 32-18

To download users and user groups on-demand:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Integration > Realms**.

Step 2 Click the download icon ( ) next to the realm where you want to download users and user groups.

What to Do Next

- Optionally, monitor the task status; see the Task Status page (**Monitoring > ASA FirePOWER Monitoring > Task Status**).

Enabling or Disabling a Realm

License: Control

Only enabled realms allow the ASA FirePOWER module to query servers. To stop queries, disable the realm.

To enable or disable a realm:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Integration > Realms**.

Step 2 Click the **State** slider next to the realm you want to enable or disable.

What to Do Next

- Optionally, monitor the task status; see the Task Status page (**Monitoring > ASA FirePOWER Monitoring > Task Status**).

Managing the Identity Policy

License: Control

To manage the Identity Policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Identity Policy**.

Step 2 If you want to copy a policy, click the copy icon ( ).

Step 3 If you want to generate a report for the policy, click the report icon ( ).

Managing Identity Rules

License: Control

To manage Identity Rules:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Identity Policy**.

Step 2 If you want to edit an identity rule, click the edit icon ( ) and make changes as described in Creating an Identity Rule, page 32-12.

Step 3 If you want to delete an identity rule, click the delete icon ( ).

Step 4 Click **Store ASA FirePOWER Changes**.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Adding an Identity Rule Category

License: Control

To add an Identity Rule Category:

Step 1 On the identity rule editor page, you have the following choices:

- Select **above Category** from the first **Insert** drop-down list, then select the category above which you want to position the rule from the second drop-down list.
- Select **below rule** from the drop-down list, then enter an existing rule number. This option is valid only when at least one rule exists in the policy.
- Select **above rule** from the drop-down list, then, enter an existing rule number. This option is valid only when at least one rule exists in the policy.

Step 2 Click **OK**.

Note Rules in a category you delete are added to the category above.

Step 3 Click **Add** or continue editing the rule.



ISE/ISE-PIC Identity Source, page 33-4

User Identity Sources

The ASA FirePOWER module supports the following identity sources:

- Authoritative *User Agent* reporting collects user data for user awareness and user access control. If you want to configure User Agents to monitor users when they log in and out of hosts or authenticate with Active Directory credentials, see The User Agent Identity Source, page 33-3.
- Authoritative *Identity Services Engine (ISE)* or *ISE-PIC* reporting collects user data for user awareness and user access control. If you have an ISE/ISE-PIC deployment and you want to configure ISE/ISE-PIC to monitor users as they authenticate via Active Directory domain controllers (DC), see The ISE/ISE-PIC Identity Source, page 33-4.
- Authoritative *captive portal authentication* actively authenticates users on your network and collects user data for user awareness and user control. If you want to configure virtual routers or Firepower Threat Defense devices to perform captive portal authentication, see The Captive Portal Active Authentication Identity Source, page 33-7.

Data from those identity sources is stored in the ASA FirePOWER module users database and the user activity database.You can configure database-server queries to automatically download new data to your module.

For more information about user detection in the ASA FirePOWER module, see User Detection Fundamentals, page 31-1.

Troubleshooting Issues with User Identity Sources

License: Any

See the following sections for information about troubleshooting issues with your identity sources.

User Agent

If you experience issues with the User Agent connection, see the *Firepower User Agent Configuration Guide*.

If you experience issues with user data reported by the User Agent, note the following:

- After the system detects activity from a User Agent user whose data is not yet in the database, the system retrieves information about them from the server. In some cases, the system requires up to 60 minutes to successfully retrieve this information from Active Directory servers. Until the data retrieval succeeds, activity seen by the User Agent user is handled by access control rules, and is not displayed in the web interface.

ISE/ISE-PIC

If you experience issues with the ISE/ISE-PIC connection, check the following:

- The pxGrid Identity Mapping feature within ISE must be enabled before you can successfully integrate ISE with the Firepower System.
- All ISE system certificates and Firepower Management Center certificates must include the **serverAuth** and **clientAuth** extended key usage values.
- The time on your ISE device must be synchronized with the time on the Firepower Management Center. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.
- If your deployment includes a primary and a secondary pxGrid node, the certificates for both nodes must be signed by the same certificate authority.
- If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.

If you experience issues with user data reported by ISE/ISE-PIC, note the following:

- After the system detects activity from an ISE user whose data is not yet in the database, the system retrieves information about them from the server. In some cases, the system requires up to 60 minutes to successfully retrieve this information from Active Directory servers. Until the data retrieval succeeds, activity seen by the ISE user is handled by access control rules, and is not displayed in the web interface.
- You cannot perform user control on ISE users who were authenticated by an LDAP, RADIUS, or RSA domain controller.
- The ASA FirePOWER module does not receive user data for ISE Guest Services users.
- Your ISE version and configuration impact how you can use ISE in the Firepower System. For more information, see The ISE/ISE-PIC Identity Source, page 33-4.
- ISE-PIC does not provide ISE attribute data.

Captive Portal

If you experience issues with captive portal authentication, note the following:

- The time on your captive portal server must be synchronized with the time on the ASA FirePOWER module.
- If you have DNS resolution configured and you create an identity rule to perform Kerberos (or HTTP Negotiate, if you want Kerberos as an option) captive portal, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN must match the hostname you provided when configuring DNS.

  For ASA with FirePOWER Services devices, the FQDN must resolve to the IP address of the routed interface used for captive portal.
- If you select Kerberos (or HTTP Negotiate, if you want Kerberos as an option) as the **Authentication Type** in an identity rule, the **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** in order to perform Kerberos captive portal active authentication.

- If you select HTTP Basic as the **Authentication Type** in an identity rule, users on your network may not notice their sessions time out. Most web browsers cache the credentials from HTTP Basic logins and use the credentials to seamlessly begin a new session after an old session times out.
- If the device you want to use for captive portal contains both inline and routed interfaces, you must configure a zone condition in your captive portal identity rules to target only the routed interfaces on the captive portal device.

The User Agent Identity Source

License: Any

The User Agent is a passive authentication method and one of the authoritative identity sources supported by the ASA FirePOWER module. When integrated with the ASA FirePOWER module, the agent monitors users when they log in and out of hosts or authenticate with Active Directory credentials. The User Agent does not report failed login attempts. The data gained from the User Agent can be used for user awareness and user control. You invoke passive authentication in your identity policy.

Installing and using User Agents allows you to perform user control; the agents associate users with IP addresses, which allows access control rules with user conditions to trigger. You can use one agent to monitor user activity on up to five Active Directory servers.

The User Agent requires a multi-step configuration, and includes the following:

- Computers or servers with the agent installed.
- Connections between an ASA FirePOWER module and the computers or Active Directory servers with the agent installed.
- Connections between the ASA FirePOWER module and the monitored LDAP servers, configured as directories within identity realms.

For detailed information about the multi-step User Agent configuration and a complete discussion of the server requirements, see the *User Agent Configuration Guide*.

The ASA FirePOWER module connection not only allows you to retrieve metadata for the users whose logins and logoffs were detected by User Agents, but also is used to specify the users and groups you want to use in access control rules. If the agent is configured to exclude specific user names, login data for those user names are not reported to the ASA FirePOWER module. User agent data is stored in the user database and user activity database on the device.

Note User Agents cannot transmit Active Directory user names ending with the $ character to the ASA FirePOWER module. You must remove the final $ character if you want to monitor these users.

If multiple users are logged into a host using remote sessions, the agent may not detect logins from that host properly. For information about how to prevent this, see the *User Agent Configuration Guide*.

Configuring a User Agent Connection

License: Control

Before you Begin

- If you plan to implement user access control, configure and enable an Active Directory realm for your User Agent connection as described in Creating a Realm, page 32-4

To configure a User Agent Connection:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Integration > Identity Sources**.

Step 2 Select **User Agent** for the **Service Type** to enable the User Agent connection.

Note To disable the connection, select **None**.

Step 3 Click the **Add New Agent** button to add a new agent.

Step 4 Type the **Hostname** or **Address** of the computer where you plan to install the agent. You must use an IPv4 address; you cannot configure the ASA FirePOWER module to connect to a User Agent using an IPv6 address.

Step 5 Click **Add**.

Step 6 To delete a connection, click the delete icon ( ) and confirm that you want to delete it.

What to Do Next

- Continue User Agent setup as described in the *Firepower User Agent Configuration Guide*.

The ISE/ISE-PIC Identity Source

License: Any

You can integrate your Cisco Identity Services Engine (ISE) or ISE Passive Identity Connector (ISE-PIC) deployment with the ASA FirePOWER module to use ISE/ISE-PIC for passive authentication. You invoke passive authentication in your identity policy.

ISE/ISE-PIC is an authoritative identity source, and provides user awareness data for users who authenticate via Active Directory (AD), LDAP, RADIUS, or RSA. Additionally, you can perform user control on AD users. ISE/ISE-PIC does not report failed login attempts or the activity of ISE Guest Services users.

Note The ASA FirePOWER module does not support 802.1x machine authentication alongside AD authentication, because the system does not associate machine authentication with users. If you use 802.1x active logins, configure ISE to report only 802.1x active logins (both machine and user). That way, a machine login is reported only once to the system.

For more information on Cisco ISE/ISE-PIC, see the *Cisco Identity Services Engine Administrator Guide* and the *Identity Services Engine Passive Identity Connector (ISE-PIC) Installation and Administrator Guide*.

Your ISE/ISE-PIC version and configuration affects its integration and interaction with the ASA FirePOWER module, as follows:

- Synchronize the time on the ISE/ISE-PIC server and the ASA FirePOWER module. Otherwise, the system may perform user timeouts at unexpected intervals.
- If you configure ISE/ISE-PIC to monitor a large number of user groups, the system may drop user mappings based on groups, due to memory limitations. As a result, access control rules with realm or user conditions may not fire as expected.

- Version 2.0 patch 4 of ISE includes support for IPv6-enabled endpoints.
- ISE-PIC does not provide ISE attribute data.

For the specific versions of ISE/ISE-PIC that are compatible with this version of the ASA FirePOWER module, see the *Cisco Firepower Compatibility Guide*.

Configuring an ISE connection populates the ASA FirePOWER module database with ISE attribute data. You can use the following ISE attributes for user awareness and user control. This is not supported with ISE-PIC.

Security Group Tags (SGT)

The Security Group Tag (SGT) specifies the privileges of a traffic source within a trusted network. Security Group Access (a feature of both Cisco TrustSec and Cisco ISE) automatically generates the SGT when a user adds a security group in TrustSec or ISE. SGA then applies the SGT attribute as packets enter the network. You can use SGTs for access control by configuring ISE as an identity source or creating custom SGT objects. For more information, see ISE SGT v. Custom SGT Rule Conditions, page 10-1.

SGT ISE attribute rule conditions can be configured in policies with or without an associated identity policy.

Endpoint Location (also known as the Location IP)

The Endpoint Location attribute is applied by Cisco ISE and identifies the IP address of the endpoint device.

You can only configure Location IP as an ISE attribute rule condition in policies with an associated identity policy.

Endpoint Profile (also known as the Device Type)

The Endpoint Profile attribute is applied by Cisco ISE and identifies the endpoint device type for each packet.

You can only configure Device Type as an ISE attribute rule condition in policies with an associated identity policy.

ISE/ISE-PIC Fields

The following fields are used to configure a connection to ISE/ISE-PIC.

Primary and Secondary Host Name/IP Address

The hostname or IP address for the primary and, optionally, the secondary ISE servers.

pxGrid Server CA

The certificate authority for the pxGrid framework. If your deployment includes a primary and a secondary pxGrid node, the certificates for both nodes must be signed by the same certificate authority.

MNT Server CA

The certificate authority for the ISE certificate when performing bulk downloads. If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.

MC Server Certificate

The certificate and key that the ASA FirePOWER module should provide to ISE when connecting to ISE or performing bulk downloads.

The **MC Server Certificate** must include the **clientAuth** extended key usage value, or it must not include any extended key usage values.

ISE Network Filter

An optional filter you can set to restrict the networks monitored by ISE. If you provide a filter, ISE monitors the networks within that filter. You can specify a filter in the following ways:

- Leave the field blank to specify any.
- Enter a single IPv4 address block using CIDR notation.
- Enter a list of IPv4 address blocks using CIDR notation, separated by commas.

Note This version of the Firepower System does not support filtering using IPv6 addresses, regardless of your ISE version.

Configuring an ISE/ISE-PIC Connection

License: Control

Before You Begin

- Configure a realm as described in Creating a Realm, page 32-4. A user download (automatic or on-demand) must be performed before you can configure an ISE attribute condition in an access control rule.

Note Configuring a realm is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions.

To configure an ISE/ISE-PIC Connection:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Integration** > **Identity Sources**.

Step 2 Select **Identity Services Engine** for the **Service Type** to enable the ISE/ISE-PIC connection.

Note To disable the connection, select **None**.

Step 3 Type a **Primary Host Name/IP Address** and, optionally, a **Secondary Host Name/IP Address**.

Step 4 Select the appropriate certificates from the **pxGrid Server CA**, **MNT Server CA**, and **MC Server Certificate** drop-down lists. Optionally, click the add icon ( ) to create an object on the fly.

Step 5 Optionally, type an **ISE Network Filter** using CIDR block notation.

Step 6 If you want to test the connection, click **Test**.

The Captive Portal Active Authentication Identity Source

License: Any

Captive portal is one of the authoritative identity sources supported by the ASA FirePOWER module. It is the only active authentication method supported by the ASA FirePOWER module, where users can authenticate onto the network through a device.

Active authentication via captive portal is performed on HTTP and HTTPS traffic only. If you want to perform captive portal on HTTPS traffic, you must create SSL rules to decrypt the traffic originating from the users you want to authenticate using captive portal.

When configured and deployed, users from specified realms authenticate through ASA FirePOWER devices in routed mode running Version 9.5(2) or later. The authentication data gained from captive portal can be used for user awareness and user control.

Captive portal also records failed authentication attempts. A failed attempt does not add a new user to the list of users in the database. The user activity type for failed authentication activity reported by captive portal is **Failed Auth User**.

You use the `captive-portal` ASA CLI command to enable captive portal for active authentication as described in the *ASA Firewall Configuration Guide* for your version: http://www.cisco.com/c/en/us/support/security/asa-5500-series-next-generation-firewalls/products-installation-and-configuration-guides-list.html. You continue configuring captive portal in your identity policy and invoke it (active authentication) in your identity rules. Identity policies are invoked in your access control policies. For more information, see Configuring Captive Portal (Active Authentication), page 32-10

Captive portal can only be performed by a device with one or more routed interfaces configured.

Note the following access control rule and SSL rule requirements:

- You must create an access control rule to allow traffic destined for the IP address and port you plan to use for captive portal. Traffic cannot be authenticated using captive portal if the destination port is not allowed in your access control policy.
- If you want to perform active authentication via captive portal on HTTPS traffic, you must create SSL rules to decrypt the traffic originating from the users you want to authenticate using captive portal.
- If you want to decrypt traffic in the captive portal connection, you must create an SSL rule to decrypt the traffic destined for the port you plan to use for captive portal.

ASA FirePOWER Module-Server Downloads

License: Any

Connections between the ASA FirePOWER module and your LDAP or AD servers allow you to retrieve user and user group metadata for certain detected users:

- LDAP and AD users authenticated by captive portal or reported by a User Agent or ISE/ISE-PIC. This metadata can be used for user awareness and user control.
- POP3 and IMAP user logins detected by traffic-based detection, if those users have the same email address as an LDAP or AD user. This metadata can be used for user awareness.

You configure an ASA FirePOWER module user database-server connection as a directory within a realm. You must select the **Download users and user groups for access control** check box to download a realm's user and user group data for user awareness and user control.

The ASA FirePOWER module obtains the following information and metadata about each user:

- LDAP user name
- first and last names
- email address
- department
- telephone number



DNS Policies

The following topics explain DNS policies, DNS rules, and how to deploy DNS policies.

- DNS Policy Overview, page 34-1
- DNS Policy Components, page 34-1
- DNS Rules, page 34-2
- DNS Policy Deploy, page 34-8

DNS Policy Overview

License: Any

DNS-based Security Intelligence allows you to whitelist or blacklist traffic based on the domain name requested by a client. Cisco provides domain name intelligence you can use to filter your traffic; you can also configure custom lists and feeds of domain names tailored to your deployment. DNS-based Security Intelligence filtering takes place after hardware-level handling and traffic decryption, and before most other policy-based inspection, analysis, or traffic handling.

Traffic blacklisted by a DNS policy is immediately blocked and therefore is not subject to any further inspection—not for intrusions, exploits, malware, and so on. You can override blacklisting with whitelisting to force access control rule evaluation, and, recommended in passive deployments, you can use a “monitor-only” setting for Security Intelligence filtering. This allows the ASA FirePOWER module to analyze connections that would have been blacklisted, but also logs the match to the blacklist and generates an end-of-connection security intelligence event.

You configure DNS-based Security Intelligence using a DNS policy and associated DNS rules. To deploy it, you must associate your DNS policy with an access control policy, then deploy your configuration.

DNS Policy Components

License: Any

A DNS policy allows you to whitelist or blacklist domain name-based connections. The following list describes the configurations you can change after creating a DNS policy.

Name and Description

Each DNS policy must have a unique name. A description is optional.

Rules

Rules provide a granular method of handling network traffic based on the domain name. Rules in a DNS policy are numbered, starting at 1. The ASA FirePOWER module matches traffic to DNS rules in top-down order by ascending rule number.

When you create a DNS policy, the ASA FirePOWER module populates it with a default Global DNS Whitelist rule, and a default Global DNS Blacklist rule. Each rule is fixed to the first position in their respective categories. You cannot modify these rules, but you can disable them. The module evaluates rules in the following order:

- Global DNS Whitelist rule (if enabled)
- whitelist rules
- Global DNS Blacklist rule (if enabled)
- blacklist and monitor rules

Usually, the module handles domain name-based network traffic according to the first DNS rule where all the rule's conditions match the traffic. If no DNS rules match the traffic, the module continues evaluating the traffic based on the associated access control policy's rules. DNS rule conditions can be simple or complex.

Editing a DNS Policy

License: Protection

Only one person should edit a DNS policy at a time, using a single browser window. If multiple users attempt to save the same policy, only the first set of saved changes are retained.

To protect the privacy of your session, after thirty minutes of inactivity on the policy editor, a warning appears. After sixty minutes, the module discards your changes.

To edit a DNS policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > DNS Policy**.

Step 2 Edit your DNS policy:

- Name and Description - To change the name or description, click the field and type the new information.
- Rules - To add, categorize, enable, disable, or otherwise manage DNS rules, click the **Rules** tab and proceed as described in Creating and Editing DNS Rules, page 34-3.

Step 3 Click **Store ASA FirePOWER Changes.**

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

DNS Rules

License: Any

DNS rules handle traffic based on the domain name requested by a host. As part of Security Intelligence, this evaluation happens after any traffic decryption, and before access control evaluation.

The ASA FirePOWER module matches traffic to DNS rules in the order you specify. In most cases, the module handles network traffic according to the first DNS rule where all the rule's conditions match the traffic. When you create DNS rules, the module places whitelist rules before monitor and blacklist rules, and evaluates traffic against whitelist rules first.

In addition to its unique name, each DNS rule has the following basic components:

State

By default, rules are enabled. If you disable a rule, the ASA FirePOWER module does not use it to evaluate network traffic, and stops generating warnings and errors for that rule.

Position

Rules in a DNS policy are numbered, starting at 1. The ASA FirePOWER module matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Conditions

Conditions specify the specific traffic the rule handles. A DNS rule must contain a DNS feed or list condition, and can also match traffic by security zone or network.

Action

A rule's action determines how the ASA FirePOWER module handles matching traffic:

- Whitelisted traffic is allowed, subject to further access control inspection.
- Monitored traffic is subject to further evaluation by remaining DNS blacklist rules. If the traffic does not match a DNS blacklist rule, it is inspected with access control rules. The module logs a Security Intelligence event for the traffic.
- Blacklisted traffic is dropped without further inspection. You can also return a Domain Not Found response, or redirect the DNS query to a sinkhole server.

Creating and Editing DNS Rules

License: Protection

In a DNS policy, you can add up to a total of 32767 DNS lists to the whitelist and blacklist rules. That is, the number of lists in the DNS policy cannot exceed 32767.

To create and edit DNS Rules:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > DNS Policy**.

Step 2 You have the following options:

- To add a new rule, click **Add DNS Rule**.
- To edit an existing rule, click the edit icon ( ).

Step 3 Enter a **Name**.

Step 4 Configure the rule components, or accept the defaults:

- Action - Select a rule **Action**; see DNS Rule Actions, page 34-5.

- Conditions - Configure the rule's conditions; see DNS Rule Conditions, page 34-6.
- Enabled - Specify whether the rule is **Enabled**.

Step 5 Click **Add** or **OK**.

Step 6 Click **Store ASA FirePOWER Changes**.

DNS Rule Management

License: Any

The Rules tab of the DNS policy editor allows you to add, edit, move, enable, disable, delete, and otherwise manage DNS rules within your policy.

For each rule, the policy editor displays its name, a summary of its conditions, and the rule action. Other icons represent warnings ( ), errors ( ), and other important information ( ). Disabled rules are dimmed and marked (`disabled`) beneath the rule name.

Enabling and Disabling DNS Rules

License: Protection

When you create a DNS rule, it is enabled by default. If you disable a rule, the ASA FirePOWER module does not use it to evaluate network traffic and stops generating warnings and errors for that rule. When viewing the list of rules in a DNS policy, disabled rules are dimmed, although you can still modify them. Note that you can also enable or disable a DNS rule using the DNS rule editor.

To enable and disable DNS Rules:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > DNS Policy**.

Step 2 In the DNS policy editor that contains the rule you want to enable or disable, right-click the rule and choose a rule state.

Step 3 Click **OK**.

Step 4 Click **Store ASA FirePOWER Changes**.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

DNS Rule Order Evaluation

License: Any

Rules in a DNS policy are numbered, starting at 1. The ASA FirePOWER module matches traffic to DNS rules in top-down order by ascending rule number. In most cases, the module handles network traffic according to the first DNS rule where all the rule's conditions match the traffic:

- For Monitor rules, the module logs the traffic, then continues evaluating traffic against lower-priority DNS blacklist rules.

- For non-Monitor rules, the module does not continue to evaluate traffic against additional, lower-priority DNS rules after that traffic matches a rule.

Note the following regarding rule order:

- The Global Whitelist is always first, and takes precedence over all other rules.
- The Whitelist section precedes the Blacklist section; whitelist rules always take precedence over other rules.
- The Global Blacklist is always first in the Blacklist section, and takes precedence over all other Monitor and blacklist rules.
- The Blacklist section contains Monitor and blacklist rules.
- When you first create a DNS rule, the module positions it last in the Whitelist section if you assign a **Whitelist** action, or last in the Blacklist section if you assign any other action.

You can drag and drop rules to reorder them, and change the evaluation order.

DNS Rule Actions

License: Any

Every DNS rule has an *action* that determines the following for matching traffic:

- handling—foremost, the rule action governs whether the module will whitelist, monitor, or blacklist traffic that matches the rule's conditions
- logging—the rule action determines when and how you can log details about matching traffic

Keep in mind that only devices deployed inline can blacklist traffic. Devices deployed passively can whitelist and log, but not affect, traffic.

Whitelist Action

The **Whitelist** action allows matching traffic to pass. When you whitelist traffic, it is subject to further inspection either by a matching access control rule, or the access control policy's default action.

The module does not log whitelist matches. However, logging of whitelisted connections depends on their eventual disposition.

Monitor Action

The **Monitor** action does not affect traffic flow; matching traffic is neither immediately whitelisted nor blacklisted. Rather, traffic is matched against additional rules to determine whether to permit or deny it. The first non-Monitor DNS rule matched determines whether the module blacklists the traffic. If there are no additional matching rules, the traffic is subject to access control evaluation.

For connections monitored by a DNS policy, the ASA FirePOWER module logs end-of-connection Security Intelligence and connection events.

Blacklist Actions

The blacklist actions blacklist traffic without further inspection of any kind:

- The **Drop** action drops the traffic.
- The **Domain Not Found** action returns a non-existent internet domain response to the DNS query, which prevents the client from resolving the DNS request.
- The **Sinkhole** action returns a sinkhole object's IPv4 or IPv6 address in response to the DNS query. The sinkhole server can log, or log and block, follow-on connections to the IP address. If you configure a **Sinkhole** action, you must also configure a sinkhole object.

For a connection blacklisted based on the **Drop** or **Domain Not Found** actions, the module logs beginning-of-connection Security Intelligence and connection events. Because blacklisted traffic is immediately denied without further inspection, there is no unique end of connection to log.

For a connection blacklisted based on the **Sinkhole** action, logging depends on the sinkhole object configuration. If you configure your sinkhole object to only log sinkhole connections, the module logs end-of-connection connection events for the follow-on connection. If you configure your sinkhole object to log and block sinkhole connections, the module logs beginning-of-connection connection events for the follow-on connection, then blocks that connection.

DNS Rule Conditions

License: Any

A DNS rule's conditions identify the type of traffic that rule handles. Conditions can be simple or complex. You must define a DNS feed or list condition. You can additionally control traffic by security zone or network.

When adding conditions to a DNS rule:

- If you do not configure a particular condition for a rule, the module does not match traffic based on that criterion.
- You can configure multiple conditions per rule. Traffic must match **all** the conditions in the rule for the rule to apply to traffic.
- For each condition in a rule, you can add up to 50 criteria. Traffic that matches **any** of a condition's criteria satisfies the condition. For example, you can use a single rule to blacklist traffic based on up to 50 DNS lists and feeds.

Controlling Traffic Based on DNS and Security Zone

License: Protection

Zone conditions in DNS rules allow you to control traffic by its source and destination security zones. A security zone is a grouping of one or more interfaces. An option you choose during a device's initial setup, called its detection mode, determines how the module initially configures the device's interfaces, and whether those interfaces belong to a security zone.

To control traffic based on DNS and security zone:

Step 1 In the DNS rule editor, click the **Zones** tab.

Step 2 Find and select the zones you want to add from the **Available Zones**. To search for zones to add, click the **Search by name** prompt above the **Available Zones** list, then type a zone name. The list updates as you type to display matching zones.

Step 3 Click to select a zone, or right-click and then select **Select All**.

Step 4 Click **Add to Source**.

Tip You can also drag and drop selected zones.

Step 5 Save or continue editing the rule.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Controlling Traffic Based on DNS and Network

License: Protection

Network conditions in DNS rules allow you to control traffic by its source IP address. You can explicitly specify the source IP addresses for the traffic you want to control.

To control traffic based on DNS and network:

Step 1 In the DNS rule editor, click the **Networks** tab.

Step 2 Find and select the networks you want to add from the **Available Networks**, as follows:

- To add a network object on the fly, which you can then add to the condition, click the add icon ( ) above the **Available Networks** list and proceed as described in Working with Network Objects, page 2-3.
- To search for network objects to add, click the **Search by name or value** prompt above the **Available Networks** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.

Step 3 Click **Add to Source**.

Tip You can also drag and drop selected objects.

Step 4 Add any source IP addresses or address blocks that you want to specify manually. Click the **Enter an IP address** prompt below the **Source Networks** list; then type an IP address or address block and click **Add**.

Step 5 Save or continue editing the rule.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

Controlling Traffic Based on DNS List, Feed, or Category

License: Protection

DNS conditions in DNS rules allow you to control traffic if a DNS list, feed, or category contains the domain name requested by the client. You must define a DNS condition in a DNS rule.

Regardless of whether you add a global or custom whitelist or blacklist to a DNS condition, the ASA FirePOWER module applies the configured rule action to the traffic. For example, if you add the Global Whitelist to a rule, and configure a **Drop** action, the module blacklists all traffic that should have been whitelisted.

To control traffic based on DNS list, feed, or category:

Step 1 In the DNS rule editor, click the **DNS** tab.

Step 2 Find and select the DNS lists and feeds you want to add from the **DNS Lists and Feeds**, as follows:

- To add a DNS list or feed on the fly, which you can then add to the condition, click the add icon ( ) above the **DNS Lists and Feeds** list and proceed as described in Working with the Intelligence Feed, page 2-6
- To search for DNS lists, feeds, or categories to add, click the **Search by name or value** prompt above the **DNS Lists and Feeds** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.

Step 3 Click **Add to Rule**.

Tip You can also drag and drop selected objects.

Step 4 Save or continue editing the rule.

What to Do Next

- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.

DNS Policy Deploy

License: Any

After you finishing updating your DNS policy configuration, you must deploy it as part of an access control policy for your changes to take effect. You must do the following:

- Associate your DNS policy with an access control policy, as described in Building the Security Intelligence Whitelist and Blacklist, page 5-3.
- Deploy configuration changes; see Deploying Configuration Changes, page 4-11.



Blocking Malware and Prohibited Files

Malicious software, or *malware*, can enter your organization's network via multiple routes. To help you identify and mitigate the effects of malware, the ASA FirePOWER module's file control and advanced malware protection components can detect, track, store, analyze, and optionally block the transmission of malware and other types of files in network traffic.

You configure the system to perform malware protection and file control as part of your overall access control configuration. *File policies* that you create and associate with access control rules handle network traffic that matches the rules.

Although you can create file policies with any license, certain aspects of malware protection and file control require that you enable specific licensed capabilities on the ASA FirePOWER module, as described in the following table.

Table 35-1 License and Appliance Requirements for Intrusion and File Inspection

Feature	Description	Add this license...
intrusion prevention	detect and optionally block intrusions and exploits	Protection
file control	detect and optionally block the transmission of file types	Protection
advanced malware protection (AMP)	detect, track, and optionally block the transmission of malware	Malware

For more information, see:

- Understanding Malware Protection and File Control, page 35-1
- Understanding and Creating File Policies, page 35-4

Understanding Malware Protection and File Control

License: Protection, Malware, or Any

Using the *advanced malware protection* feature, you can configure the ASA FirePOWER module to detect, track, analyze, and optionally block malware files being transmitted on your network.

The system can detect and optionally block malware in many types of files, including PDFs, Microsoft Office documents, and others. ASA FirePOWER modules monitor specific application protocol-based network traffic for transmissions of those file types. When the ASA FirePOWER module detects an eligible file, the ASA FirePOWER module then performs a *malware cloud lookup* using the file's SHA-256 hash value. Based on these results, the Cisco cloud returns a file disposition to the ASA FirePOWER module.

If a file has a disposition in the cloud that you know to be incorrect, you can add the file's SHA-256 value to a file list:

- To treat a file as if the cloud assigned a clean disposition, add the file to the *clean list*.
- To treat a file as if the cloud assigned a malware disposition, add the file to the *custom detection list.*

If the system detects a file's SHA-256 value on a file list, it takes the appropriate action without performing a malware lookup or checking the file disposition. Note that you must configure a rule in the file policy with either a **Malware Cloud Lookup** or **Block Malware** action and a matching file type to calculate a file's SHA value. You can enable use of the clean list or custom detection list on a per-file-policy basis.

To inspect or block files, you must enable a Protection license on the ASA FirePOWER module. To add files to a file list, you must also enable a Malware license.

Understanding File Dispositions

The system determines file dispositions based on the disposition returned by the Cisco cloud. A file can have one of the following file dispositions returned by the Cisco cloud, as a result of addition to a file list, or due to threat score:

- `Malware` indicates that the cloud categorized the file as malware.
- `Clean` indicates that the cloud categorized the file as clean, or that a user added the file to the clean list.
- `Unknown` indicates that a malware cloud lookup occurred before the cloud assigned a disposition. The cloud has not categorized the file.
- `Custom Detection` indicates that a user added the file to the custom detection list.
- `Unavailable` indicates that the ASA FirePOWER module could not perform a malware cloud lookup. You may see a small percentage of events with this disposition; this is expected behavior.



Tip If you see several `Unavailable` malware events in quick succession, check your cloud connection and port configuration. For more information, see Security, Internet Access, and Communication Ports, page D-1.

Based on the file disposition, the ASA FirePOWER module either blocks the file or blocks its upload or download. To improve performance, if the system already knows the disposition for a file based on its SHA-256 value, your appliance uses the cached disposition rather than querying the Cisco cloud.

Note that file dispositions can change. For example, the cloud can determine that a file that was previously thought to be clean is now identified as malware, or the reverse—that a malware-identified file is actually clean. When the disposition changes for a file for which you performed a malware lookup in the last week, the cloud notifies the ASA FirePOWER module so the system can take appropriate action the next time it detects that file being transmitted. A changed file disposition is called a *retrospective* disposition.

File dispositions returned from a malware cloud lookup have a time-to-live (TTL) value. After a file disposition has been held for the duration specified in the TTL value without update, the system purges the cached information. Dispositions have the following TTL values:

- Clean—4 hours
- Unknown—1 hour
- Malware—1 hour

If a malware cloud lookup against the cache identifies a cached disposition that timed out, the system performs a fresh lookup to determine a file disposition.

Understanding File Control

If your organization wants to block not only the transmission of malware files, but all files of a specific type (regardless of whether the files contain malware), the *file control* feature allows you to cast a wider net. As with malware protection, the ASA FirePOWER module monitors network traffic for transmissions of specific file types, then either blocks or allows the file.

File control is supported for all file types where the system can detect malware, plus many additional file types. These file types are grouped into basic categories, including multimedia (swf, mp3), executables (exe, torrent), and PDFs. Note that file control, unlike malware protection, does not require queries of the Cisco cloud.

Configuring Malware Protection and File Control

License: Protection or Malware

You configure malware protection and file control as part of your overall access control configuration by associating file policies with access control rules. This association ensures that before the system passes a file in traffic that matches an access control rule's conditions, it first inspects the file.

A file policy, like its parent access control policy, contains rules that determine how the system handles files that match the conditions of each rule. You can configure separate file rules to take different actions for different file types, application protocols, or directions of transfer.

When a file matches a rule, the rule can:

- allow or block files based on simple file type matching
- block files based on malware file disposition

In addition, the file policy can automatically treat a file as if it is clean or malware based on entries in the clean list or custom detection list

As a simple example, you could implement a file policy that blocks your users from downloading executable files. For detailed information on file policies and associating them with access control rules, see Understanding and Creating File Policies, page 35-4.

Logging Events Based on Malware Protection and File Control

License: Protection or Malware

The ASA FirePOWER module logs records of the system's file inspection and handling file events, and malware events:

- *File events* represent files that the system detected, and optionally blocked, in network traffic.

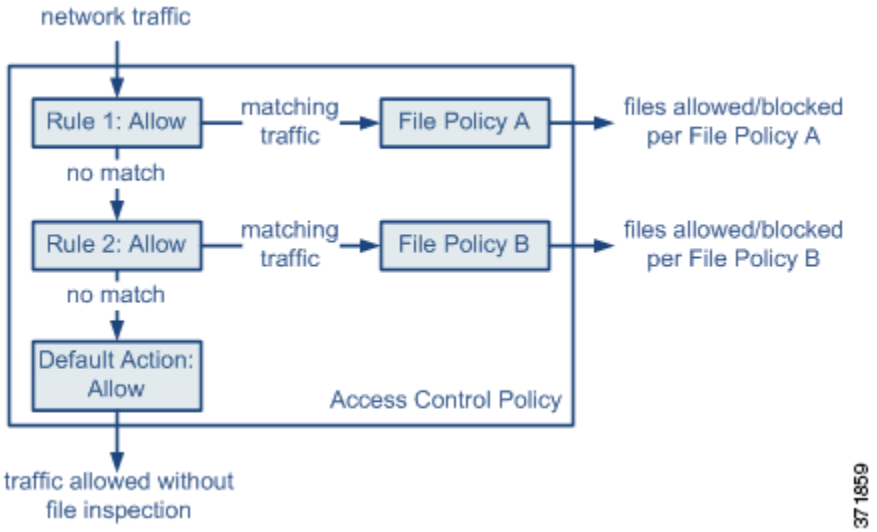
- *Malware events* represent malware files detected, and optionally blocked, in network traffic by the system.
- *Retrospective malware events* represent files whose malware file dispositions have changed.

When the system generates a malware event based on detection or blocking of malware in network traffic, it also generates a file event, because to detect malware in a file the system must first detect the file itself.

Understanding and Creating File Policies

License: Protection or Malware

A file policy is a set of configurations that the system uses to perform advanced malware protection and file control, as part of your overall access control configuration.



The policy has two access control rules, both of which use the Allow action and are associated with file policies. The policy's default action is also to allow traffic, but without file policy inspection. In this scenario, traffic is handled as follows:

- Traffic that matches `Rule 1` is inspected by `File Policy A`.
- Traffic that does not match `Rule 1` is evaluated against `Rule 2`. Traffic that matches `Rule 2` is inspected by `File Policy B`.
- Traffic that does not match either rule is allowed; you cannot associate a file policy with the default action.

A file policy, like its parent access control policy, contains rules that determine how the system handles files that match the conditions of each rule. You can configure separate file rules to take different actions for different file types, application protocols, or directions of transfer.

Once a file matches a rule, the rule can:

- allow or block files based on simple file type matching
- block files based on Malware file disposition

In addition, the file policy can automatically treat a file as if it is clean or malware based on entries in the clean list or custom detection list

You can associate a single file policy with an access control rule whose action is **Allow**, **Interactive Block**, or **Interactive Block with reset**. The system then uses that file policy to inspect network traffic that meets the conditions of the access control rule. By associating different file policies with different access control rules, you have granular control over how you identify and block files transmitted on your network. Note, however, that you **cannot** use a file policy to inspect traffic handled by the access control default action. For detailed information, see Inspecting Allowed Traffic For Intrusions and Malware, page 11-2.

File Rules

You populate a file policy with file rules. The following table describes the components of a file rule.

Table 35-2 File Rule Components

File Rule Component	Description
application protocol	The system can detect and inspect files transmitted via FTP, HTTP, SMTP, IMAP, POP3, and NetBIOS-ssn (SMB). To improve performance, you can restrict file detection to only one of those application protocols on a per-file rule basis.
direction of transfer	You can inspect incoming FTP, HTTP, IMAP, POP3, and NetBIOS-ssn (SMB) traffic for downloaded files; you can inspect outgoing FTP, HTTP, SMTP, and NetBIOS-ssn (SMB) traffic for uploaded files.
file categories and types	The system can detect various types of files. These file types are grouped into basic categories, including multimedia (swf, mp3), executables (exe, torrent), and PDFs. You can configure file rules that detect individual file types, or on entire categories of file types. For example, you could block all multimedia files, or just ShockWave Flash (swf) files. Or, you could configure the system to alert you when a user downloads a BitTorrent (torrent) file. ⚠ **Caution** Frequently triggered file rules can affect system performance. For example, detecting multimedia files in HTTP traffic (YouTube, for example, transmits significant Flash content) could generate an overwhelming number of events.
file rule action	A file rule's action determines how the system handles traffic that matches the conditions of the rule. **Note** File rules are evaluated in rule-action, not numerical, order. For more information, see the next section, File Rule Actions and Evaluation Order.

File Rule Actions and Evaluation Order

Each file rule has an associated action that determines how the system handles traffic that matches the conditions of the rule. You can set separate rules within a file policy to take different actions for different file types, application protocols, or directions of transfer. The rule actions are as follows, in rule-action order:

- *Block Files* rules allow you to block specific file types.

- *Block Malware* rules allow you to calculate the SHA-256 hash value of specific file types, then use a cloud lookup process to first determine if files traversing your network contain malware, then block files that represent threats.
- *Malware Cloud Lookup* rules allow you to log the malware disposition of files traversing your network based on a cloud lookup, while still allowing their transmission.
- *Detect Files* rules allow you to log the detection of specific file types while still allowing their transmission.

For each file rule action, you can configure options to reset the connection when a file transfer is blocked. The following table details the options available to each file action.

Table 35-3 File Rule Actions

Action	Resets Connection?
Block Files	yes (recommended)
Block Malware	yes (recommended)
Detect Files	no
Malware Cloud Lookup	no

File and Malware Detection, Capture, and Blocking Notes and Limitations

Note the following details and limitations on file and malware detection, capture, and blocking behavior:

- Until a file is detected and block in a session, packets from the session may be subject to intrusion inspection.
- If an end-of-file marker is not detected for a file, regardless of transfer protocol, the file is not blocked by a **Block Malware** rule or by the custom detection list. The system waits to block the file until the entire file has been received, as indicated by the end-of-file marker, and blocks the file after the marker is detected.
- If the end-of-file marker for an FTP file transfer is transmitted separately from the final data segment, the marker is blocked and the FTP client indicates that the file transfer failed, but the file actually completely transfers to disk.
- FTP transfers commands and data over different channels. In a passive deployment, the traffic from an FTP data session and its control session may not be load-balanced to the same Snort.
- If a file matches a rule with an application protocol condition, file event generation occurs after the system successfully identifies a file's application protocol. Unidentified files do not generate file events.
- For an access control policy using a file policy with **Block Malware** rules for FTP, if you set the default action to an intrusion policy with **Drop when Inline** disabled, the system generates events for detected files or malware matching the rules, but does not drop the files. To block FTP fire transfers and use an intrusion policy as the default action for the access control policy where you select the file policy, you must select an intrusion policy with **Drop when Inline** enabled.
- File rules with **Block Files** and **Block Malware** actions block automatic resumption of file download via HTTP by blocking new sessions with the same file, URL, server, and client application detected for 24 hours after the initial file transfer attempt occurs.
- In rare cases, if traffic from an HTTP upload session is out of order, the system cannot reassemble the traffic correctly and therefore does not block it or generate a file event.

- If you transfer a file over NetBIOS-ssn (such as an SMB file transfer) that is blocked with a **Block Files** rule, you may see a file on the destination host. However, the file is unusable because it is blocked after the download starts, resulting in an incomplete file transfer.
- If you create file rules to detect or block files transferred over NetBIOS-ssn (such as an SMB file transfer), the system does not inspect files transferred in an established TCP or SMB session started before you apply an access control policy invoking the file policy, so those files will not be detected or blocked.
- A rule configured to block files in a passive deployment does not block matching files. Because the connection continues to transmit the file, if you configure the rule to log the beginning of the connection, you may see multiple events logged for this connection.
- If the total number of bytes for all file names for files in a POP3, POP, SMTP, or IMAP session exceeds 1024, file events from the session may not reflect the correct file names for files that were detected after the file name buffer filled.
- When transmitting text-based files over SMTP, some mail clients convert newlines to the CRLF newline character standard. Since Mac-based hosts use the carriage return (CR) character and Unix/Linux-based hosts use the line feed (LF) character, newline conversion by the mail client may modify the size of the file. Note that some mail clients default to newline conversion when processing an unrecognizable file type.
- Cisco recommends that you enable **Reset Connection** for the **Block Files** and **Block Malware** actions to prevent blocked application sessions from remaining open until the TCP connection resets. If you do not reset connections, the client session remains open until the TCP connection resets itself.
- If a file rule is configured with a **Malware Cloud Lookup** or **Block Malware** action and the ASA FirePOWER module cannot establish connectivity with the cloud, the system cannot perform any configured rule action options until cloud connectivity is restored.

File Rule Evaluation Example

Unlike in access control policies, where rules are evaluated in numerical order, file policies handle files in File Rule Actions and Evaluation Order, page 35-5. That is, simple blocking takes precedence over malware inspection and blocking, which takes precedence over simple detection and logging. As an example, consider four rules that handle PDF files in a single file policy. Regardless of the order in which they appear in the module interface, these rules are evaluated in the following order:

Table 35-4 File Rule Evaluation Order Example

App. Protocol	Direction	Action	Action Options	Result
SMTP	Upload	Block Files	Reset Connection	Blocks users from emailing PDF files and resets the connection.
FTP	Download	Block Malware	Reset Connection	Blocks the download of malware PDF files via file transfer, and resets the connection.
POP3 IMAP	Download	Malware Cloud Lookup	none	Inspects PDF files received via email for malware.
Any	Any	Detect Files	none	Detects and logs, but allows the traffic, when users view PDF files on the web (that is, via HTTP).

The ASA FirePOWER module uses warning icons ( ) to designate conflicting file rules.

Note that you cannot perform malware analysis on all file types detected by the system. After you select values from the **Application Protocol**, **Direction of Transfer**, and **Action** drop-down lists, the system constrains the list of file types.

Logging File Events, Malware Events and Alerts

When you associate a file policy with an access control rule, the system automatically enables file and malware event logging for matching traffic. When the system inspects a file, it can generate the following types of events:

- *file events*, which represent detected or blocked files, as well as detected malware files
- *malware events*, which represent detected malware files
- *retrospective malware events*, which are generated when the Malware file disposition for a previously detected file changes

When a file policy generates a file or malware event, or captures a file, the system automatically logs the end of the associated connection, regardless of the logging configuration of the invoking access control rule.

Note File events generated by inspecting NetBIOS-ssn (SMB) traffic do not immediately generate connection events because the client and server establish a persistent connection. The system generates connection events after the client or server ends the session.

For each of these connection events:

- The **Files** field contains an icon ( ) that indicates the number of files (including malware files) detected in the connection; click the icon to see a list of those files and, for malware files, their file dispositions.
- The **Reason** field indicates the reason the connection event was logged, which depends on the file rule action:
  - `File Monitor` for Detect Files and Malware Cloud Lookup file rules and for files on the clean list
  - `File Block` for Block Files or Block Malware file rules
  - `File Custom Detection` if the system encountered a file on the custom detection list
  - `File Resume Allow` where file transmission was originally blocked by a Block Files or Block Malware file rule. After a new access control policy was applied that allowed the file, the HTTP session automatically resumed.
  - `File Resume Block` where file transmission was originally allowed by a Detect Files or Malware Cloud Lookup file rule. After a new access control policy was applied that blocked the file, the HTTP session automatically stopped.
- For connections where a file or malware was blocked, the **Action** is `Block`.

As with any kind of event generated by the ASA FirePOWER module, you can view file and malware events. You can also use malware events to alert you via SNMP or syslog.

Internet Access

The system uses port 443 to perform malware cloud lookups for network-based AMP. You must open that port outbound on the ASA FirePOWER module.

Managing File Policies

You create, edit, delete, and compare file policies on the File Policies page (**Policies > Files**), which displays a list of existing file policies along with their last-modified dates.

Clicking the apply icon ( ) for a file policy displays a dialog box that tells you which access control policies use the file policy, then redirects you to the Access Control Policy page. This is because you cannot apply a file policy independently, as a file policy is considered part of its parent access control policies. To use a new file policy, or to apply changes made to an existing file policy, you must apply or reapply the parent access control policies.

Note that you cannot delete a file policy used in a saved or applied access control policy.

For more information on managing file policies, see the following sections:

- Creating a File Policy, page 35-9
- Working with File Rules, page 35-10
- Comparing Two File Policies, page 35-12

Creating a File Policy

License: Protection or Malware

After you create a file policy and populate it with rules, you can use it in an access control policy.

Tip To make a copy of an existing file policy, click the copy icon ( ), then type a unique name for the new policy in the dialog box that appears. You can then modify the copy.

To create a file policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Files**.

The File Policies page appears.

Step 2 Click **New File Policy**.

The New File Policy dialog box appears.

For a new policy, the module interface indicates that the policy is not in use. If you are editing an in-use file policy, the module interface tells you how many access control policies use the file policy. In either case, you can click the text to jump to the Access Control Policies page; see Getting Started with Access Control Policies, page 4-1.

Step 3 Enter a **Name** and optional **Description** for your new policy, then click **Save**.

The File Policy Rules tab appears.

Step 4 Add one or more rules to the file policy.

File rules give you granular control over which file types you want to log, block, or scan for malware. For information on adding file rules, see Working with File Rules, page 35-10.

Step 5 Configure the advanced options. See Configuring Advanced File Policy General Options, page 35-11 for more information.

Step 6 Click **Store ASA FirePOWER Changes**.

To use your new policy, you must add the file policy to an access control rule, then apply the access control policy. If you are editing an existing file policy, you must reapply any access control policies that use the file policy.

Working with File Rules

License: Protection or Malware

To be effective, a file policy must contain one or more rules. You create, edit, and delete rules on the File Policy Rules page, which appears when you create a new file policy or edit an existing policy. The page lists all the rules in the policy, along with each rule's basic characteristics.

The page also notifies you of how many access control policies use this file policy. You can click the notification to display a list of the parent policies and, optionally, continue to the Access Control Policies page.

To create a file rule:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Files**.

The File Policies page appears.

Step 2 You have the following options:

- To add rules to a new policy, click **New File Policy** to create a new policy; see Creating a File Policy, page 35-9.
- To add rules to an existing policy, click the edit icon ( ) next to the policy.

Step 3 On the File Policy Rules page that appears, click **Add File Rule**.

The Add File Rule dialog box appears.

Step 4 Select an **Application Protocol** from the drop-down list.

Any, the default, detects files in HTTP, SMTP, IMAP, POP3, FTP, and NetBIOS-ssn (SMB) traffic.

Step 5 Select a **Direction of Transfer** from the drop-down list.

You can inspect the following types of incoming traffic for downloaded files:

- HTTP
- IMAP
- POP3
- FTP
- NetBIOS-ssn (SMB)

You can inspect the following types of outgoing traffic for uploaded files:

- HTTP
- FTP
- SMTP
- NetBIOS-ssn (SMB)

Use **Any** to detect files over multiple application protocols, regardless of whether users are sending or receiving.

Step 6 Select a file rule **Action**. See the File Rule Actions table for more information.

When you select either **Block Files** or **Block Malware**, **Reset Connection** is enabled by default. To **not** reset the connection where a blocked file transfer occurs, clear the **Reset Connection** check box.

Note Cisco recommends that you leave **Reset Connection** enabled to prevent blocked application sessions from remaining open until the TCP connection resets.

For detailed information on file rule actions, see File Rule Actions and Evaluation Order, page 35-5.

Step 7 Select one or more **File Types**. Use the Shift and Ctrl keys to select multiple file types. You can filter the list of file types in the following ways:

- Select one or more **File Type Categories**.
- Search for a file type by its name or description. For example, type `Windows` in the **Search name and description** field to display a list of Microsoft Windows-specific files.

The file types that you can use in a file rule vary depending on your selections for **Application Protocol**, **Direction of Transfer**, and **Action**.

For example, selecting **Download** as the **Direction of Transfer** removes `GIF`, `PNG`, `JPEG`, `TIFF`, and `ICO` from the **Graphics** category to prevent an excess of file events.

Step 8 Add the selected file types to the **Selected Files Categories and Types** list:

- Click **Add** to add selected file types to the rule.
- Drag and drop one or more file types into the **Selected Files Categories and Types** list.
- With a category selected, click **All types in selected Categories**, then either click **Add** or drag and drop that selection to the **Selected Files Categories and Types** list.

Step 9 Click **Store ASA FirePOWER Changes**.

The file rule is added to the policy. If you are editing an existing file policy, you must reapply any access control policies that use the file policy for your changes to take effect.

Configuring Advanced File Policy General Options

License: Malware

In a file policy, you can set the following advanced options in the General section.

Table 35-5 Advanced File Policy General Options

Field	Description	Default Value
Enable Custom Detection List	Select this to block files on the custom detection list when detected.	enabled
Enable Clean List	Select this to allow files on the clean list when detected.	enabled

To configure advanced file policy general options:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Files**.

The File Policies page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

The File Policy Rule page appears.

Step 3 Select the **Advanced** tab.

The Advanced tab appears.

Step 4 Modify the options as described in the Advanced File Policy General Options table.

Step 5 Click **Store ASA FirePOWER Changes**.

You must reapply any access control policies that use the file policy you edited.

Comparing Two File Policies

License: Protection

To review policy changes for compliance with your organization's standards or to optimize system performance, you can examine the differences between any two file policies, or two revisions of the same policy.

The file policy *comparison view* displays two file policies or revisions in a side-by-side format, with the time of last modification and the last user to modify displayed next to each policy name. Differences between the two policies are highlighted:

- Blue indicates that the highlighted setting is different in the two policies, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one policy but not the other.

You can navigate through the differences by clicking **Previous** and **Next**. The double-arrow icon ( ) centered between the left and right sides moves, and the **Difference** number adjusts to identify which difference you are viewing. Optionally, you can generate a file policy *comparison report*, which is a PDF version of the comparison view.

To compare two file policies:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Files.**

The File Policies page appears.

Step 2 Click **Compare Policies**.

The Select Comparison dialog box appears.

Step 3 From the **Compare Against** drop-down list, select the type of comparison you want to make:

- To compare two different policies, select either **Running Configuration** or **Other Policy**. The practical difference between the two options is that if you select **Running Configuration**, the system limits one of your comparison choices to the set of currently applied file policies.
- To compare revisions of the same policy, select **Other Revision**.

The dialog box refreshes, displaying your comparison options.

Step 4 Depending on the comparison type you selected, you have the following choices:

- If you are comparing two different policies, select the policies you want to compare: **Policy A** or **Target/Running Configuration A**, and **Policy B**.

- If you are comparing revisions of the same policy, select the **Policy** you want to use, then select the two revisions: **Revision A** and **Revision B**. Revisions are listed by date and user name.

Step 5 Click **OK**.

The comparison view appears.

Step 6 Optionally, click **Comparison Report** to generate a file policy comparison report. You are prompted to save the report to your computer.



Logging Connections in Network Traffic

As devices monitor traffic generated by the hosts on your network, they can generate logs of the connections they detect. Various settings in access control and SSL policies give you granular control over which connections you log, when you log them, and where you store the data. An access control rule's specific logging configuration also determines whether you log file and malware events associated with the connection.

In most cases, you can log a connection at its beginning and its end. When you log a connection, the system generates a *connection event*. You can also log a special kind of connection event, called a *Security Intelligence event*, whenever a connection is blacklisted (blocked) by the reputation-based Security Intelligence feature.

Connection events contain data about the detected sessions.

You should log connections according to the security and compliance needs of your organization.

For more information on logging connection data, see:

- Deciding Which Connections To Log, page 36-1
- Logging Security Intelligence (Blacklisting) Decisions, page 36-8
- Logging Connections Based on Access Control Handling, page 36-9
- Logging URLs Detected in Connections, page 36-13
- Logging Encrypted Connections, page 36-14

Deciding Which Connections To Log

License: Any

Using various settings in access control and SSL policies, you can log any connection that your ASA FirePOWER module monitors. In most cases, you can log a connection at its beginning and its end. However, because blocked traffic is immediately denied without further inspection, the system can log only beginning-of-connection events for blocked or blacklisted traffic; there is no unique end of connection to log.

When you log a connection event, you can view it in the event viewer. You can also send connection data to an external syslog or SNMP trap server.

Tip To perform detailed analysis of connection data using the ASA FirePOWER module, Cisco recommends you log the ends of critical connections.

For more information, see:

- Logging Critical Connections, page 36-2
- Logging the Beginning and End of Connections, page 36-3
- Logging Connections to the ASA FirePOWER Module or External Server, page 36-4
- Understanding How Access Control and SSL Rule Actions Affect Logging, page 36-4
- License Requirements for Connection Logging, page 36-7

Logging Critical Connections

License: Any

You should log connections according to the security and compliance needs of your organization. If your goal is to limit the number of events you generate and improve performance, only enable logging for the connections critical to your analysis. However, if you want a broad view of your network traffic for profiling purposes, you can enable logging for additional connections. Various settings in access control and SSL policies give you granular control over which connections you log, when you log them, and where you store the data.

Caution Logging blocked TCP connections during a Denial of Service (DoS) attack can overwhelm the system with multiple similar events. Before you enable logging for an Block rule, consider whether the rule monitors traffic on an Internet-facing interface or other interface vulnerable to DoS attack.

In addition to the logging that you configure, the system automatically logs most connections where the system detects a prohibited file, malware, or intrusion attempt. The system saves these end-of-connection events for further analysis. All connection events reflect why they were automatically logged using the Action and Reason fields.

Security Intelligence Blacklisting Decisions (Optional)

You can log a connection whenever it is blacklisted (blocked) by the reputation-based Security Intelligence feature. Optionally, and recommended in passive deployments, you can use a monitor-only setting for Security Intelligence filtering. This allows the system to further analyze connections that would have been blacklisted, but still log the match to the blacklist.

When you enable Security Intelligence logging, blacklist matches generate Security Intelligence events as well as connection events. A Security Intelligence event is a special kind of connection event that you can view and analyze separately, and that is also stored and pruned separately. For more information, see Logging Security Intelligence (Blacklisting) Decisions, page 36-8.

Access Control Handling (Optional)

You can log a connection when it is handled by an access control rule or the access control default action. You configure this logging on a per-access control rule basis so that you only log critical connections. For more information, see Logging Connections Based on Access Control Handling, page 36-9.

Connections Associated with Intrusions (Automatic)

When an intrusion policy invoked by an access control rule (see Tuning Traffic Flow Using Access Control Rules, page 6-1) detects an intrusion and generates an intrusion event, the system automatically logs the end of the connection where the intrusion occurred, regardless of the logging configuration of the rule.

However, when an intrusion policy associated with the access control default action (see Setting Default Handling and Inspection for Network Traffic, page 4-4) generates an intrusion event, the system does **not** automatically log the end of the associated connection. Instead, you must explicitly enable default action connection logging. This is useful for intrusion prevention-only deployments where you do not want to log any connection data.

For connections where an intrusion was blocked, the action for the connection in the connection log is `Block`, with a reason of `Intrusion Block`, even though to perform intrusion inspection you must use an Allow rule.

Connections Associated with File and Malware Events (Automatic)

When a file policy invoked by an access control rule detects a prohibited file (including malware) and generates a file or malware event, the system automatically logs the end of the connection where the file was detected, regardless of the logging configuration of the access control rule. You **cannot** disable this logging.

Note File events generated by inspecting NetBIOS-ssn (SMB) traffic do not immediately generate connection events because the client and server establish a persistent connection. The system generates connection events after the client or server ends the session.

For connections where a file was blocked, the action for the connection in the connection log is `Block` even though to perform file and malware inspection you must use an Allow rule. The connection's reason is either `File Monitor` (a file type or malware was detected), or `Malware Block` or `File Block` (a file was blocked).

Logging the Beginning and End of Connections

License: Any

When the system detects a connection, in most cases you can log it at its beginning and its end.

However, because blocked traffic is immediately denied without further inspection, in most cases you can log only beginning-of-connection events for blocked or blacklisted traffic; there is no unique end of connection to log.

Note For a single non-blocked connection, the end-of-connection event contains all of the information in the beginning-of-connection event, as well as information gathered over the duration of the session.

Note that monitoring a connection for any reason forces end-of-connection logging; see Understanding Logging for Monitored Connections, page 36-5.

The following table details the differences between beginning and end-of-connection events, including the advantages to logging each.

Table 36-1 Comparing Beginning and End-of-Connection Events

	Beginning-of-Connection Events	End-of-Connection Events
Can be generated...	when the system detects the beginning of a connection (or, after the first few packets if event generation depends on application or URL identification)	when the system: • detects the close of a connection • does not detect the end of a connection after a period of time • can no longer track the session due to memory constraints
Can be logged for...	all connections evaluated by Security Intelligence or access control rules	all connections are configurable, though the system cannot log the end of blocked or blacklisted connections
Contain...	only information that can be determined in the first packet (or the first few packets, if event generation depends on application or URL identification)	all information in the beginning-of-connection event, plus information determined by examining traffic over the duration of the session, for example, the total amount of data transmitted or the timestamp of the last packet in the connection
Are useful...	if you want to log: • blocked connections, including Security Intelligence blacklisting decisions	if you want to: • perform any kind of detailed analysis on information collected over the duration of the session • view connection data in graphical format

Logging Connections to the ASA FirePOWER Module or External Server

License: Any

You can log connection events to the ASA FirePOWER module, as well as to an external syslog or SNMP trap server. Before you can log connection data to an external server, you must configure a connection to that server called an *alert response*; see Working with Alert Responses, page 38-2.

Understanding How Access Control and SSL Rule Actions Affect Logging

License: feature dependent

Every access control and SSL rule has an *action* that determines not only how the system inspects and handles the traffic that matches the rule, but also when and how you can log details about matching traffic.

For more information, see:

- Using Rule Actions to Determine Traffic Handling and Inspection, page 6-7
- Understanding Logging for Monitored Connections, page 36-5
- Understanding Logging for Trusted Connections, page 36-5
- Understanding Logging for Blocked and Interactively Blocked Connections, page 36-5
- Understanding Logging for Allowed Connections, page 36-6

- Disabling File and Malware Event Logging for Allowed Connections, page 36-7

Understanding Logging for Monitored Connections

License: feature dependent

The system always logs the ends of the following connections to the ASA FirePOWER module, regardless of the logging configuration of the rule or default action that later handles the connection:

- connections matching a Security Intelligence blacklist set to monitor
- connections matching an access control Monitor rule

In other words, if a packet matches a Monitor rule or Security Intelligence monitored blacklist, the connection is always logged, even if the packet matches no other rules and you do not enable logging on the default action. Whenever the system logs a connection event as the result of Security Intelligence filtering, it also logs a matching Security Intelligence event, which is a special kind of connection event that you can view and analyze separately; see Logging Security Intelligence (Blacklisting) Decisions, page 36-8.

Because monitored traffic is always later handled by another rule or by the default action, the action associated with a connection logged due to a monitor rule is never `Monitor`. Rather, it reflects the action of the rule or default action that later handles the connection.

The system does **not** generate a separate event each time a single connection matches an SSL or access control Monitor rule. Because a single connection can match multiple Monitor rules, each connection event logged to the ASA FirePOWER module can include and display information on the first eight Monitor access control rules that the connection matches, as well as the first matching Monitor SSL rule.

Similarly, if you send connection events to an external syslog or SNMP trap server, the system does not send a separate alert each time a single connection matches a Monitor rule. Rather, the alert that the system sends at the end of the connection contains information on the Monitor rules the connection matched.

Understanding Logging for Trusted Connections

License: feature dependent

A trusted connection is one that is handled by a Trust access control rule or the default action in an access control policy. You can log the beginnings and ends of these connections; however, keep in mind that trusted connections, regardless of whether they are encrypted, are not inspected for intrusions, or prohibited files and malware. Therefore, connection events for trusted connections contain limited information.

Understanding Logging for Blocked and Interactively Blocked Connections

License: feature dependent

For access control rules and access control policy default actions that block traffic (including interactive blocking rules), the system logs **beginning**-of-connection events. Matching traffic is denied without further inspection.

Connection events for sessions blocked by an access control or SSL rule have an action of `Block` or `Block with reset`. Blocked encrypted connections have a reason of `SSL Block`.

Interactive blocking access control rules, which cause the system to display a warning page when a user browses to a prohibited website, log ends of connections. This is because if the user clicks through the warning page, the connection is considered a new, allowed connection which the system can monitor and

log; see Understanding Logging for Allowed Connections, page 36-6.

Therefore, for packets that match an Interactive Block or Interactive Block with reset rule, the system can generate the following connection events:

- a beginning-of-connection event when a user's request is initially blocked and the warning page is displayed; this event has an associated action of `Interactive Block` or `Interactive Block with reset`
- multiple beginning- or end-of-connection events if the user clicks through the warning page and loads the originally requested page; these events have an associated action of `Allow` and a reason of `User Bypass`

Note that only devices deployed inline can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.

Caution Logging blocked TCP connections during a Denial of Service (DoS) attack can overwhelm the system with multiple similar events. Before you enable logging for an Block rule, consider whether the rule monitors traffic on an Internet-facing interface or other interface vulnerable to DoS attack.

Understanding Logging for Allowed Connections

License: feature dependent

Decrypt SSL rules, Do not decrypt SSL rules, and Allow access control rules permit matching traffic to pass to the next phase of inspection and traffic handling.

When you allow traffic with an access control rule, you can use an associated intrusion or file policy (or both) to further inspect traffic and block intrusions, prohibited files, and malware before the traffic can reach its final destination.

Connections for traffic matching an Allow access control rule are logged as follows:

- When an intrusion policy invoked by an access control rule detects an intrusion and generates an intrusion event, the system automatically logs the end of the connection where the intrusion occurred to the ASA FirePOWER module, regardless of the logging configuration of the rule.
- When a file policy invoked by an access control rule detects a prohibited file (including malware) and generates a file or malware event, the system automatically logs the end of the connection where the file was detected to the ASA FirePOWER module, regardless of the logging configuration of the access control rule.
- Optionally, you can enable beginning- and end-of-connection logging for any allowed traffic, including traffic that the system deems safe or that you do not inspect with an intrusion or file policy.

For all of the resulting connection events, the Action and Reason fields reflect why the events were logged. Note that:

- An action of `Allow` represents explicitly allowed and user-bypassed interactively blocked connections that reached their final destination.
- An action of `Block` represents a connection that was at first allowed by an access control rule, but where an intrusion, prohibited file, or malware was detected.

Disabling File and Malware Event Logging for Allowed Connections

License: Protection or Malware

When you allow unencrypted or decrypted traffic with an access control rule, you can use an associated file policy to inspect transmitted files, and block prohibited files and malware before it can reach its destination; see Tuning Intrusion Prevention Performance, page 11-6.

When the system detects a prohibited file, it automatically logs one of the following types of event to the ASA FirePOWER module:

- *file events*, which represent detected or blocked files, including malware files
- *malware events*, which represent detected or blocked malware files only
- *retrospective malware events*, which are generated when the malware disposition for a previously detected file changes

If you do not want to log file or malware events, you can disable this logging on a per-access-control-rule basis by clearing the **Log Files** check box on the Logging tab of the access control rule editor.

Note Cisco recommends you leave file and malware event logging enabled.

Regardless of whether you save file and malware events, when network traffic violates a file policy, the system automatically logs the end of the associated connection to the ASA FirePOWER module, regardless of the logging configuration of the invoking access control rule; see Connections Associated with File and Malware Events (Automatic), page 36-3.

License Requirements for Connection Logging

License: feature dependent

Because you configure connection logging in access control and SSL policies, you can log any connection that those policies can successfully handle.

Although you can create access control and SSL policies regardless of the licenses on your ASA FirePOWER module, certain aspects of access control require that you enable specific licensed capabilities before you can apply the policy.

The following table explains the licenses you must have to successfully configure access control, and therefore to log connections handled by an access control policy.

Table 36-2 License Requirements for Connection Logging in Access Control Policies

To log connections...	License
for traffic handled using, network, port, or literal URL criteria	Any
for traffic handled using geolocation data	Any
associated with: • IP addresses with a poor reputation (Security Intelligence filtering) • intrusions or prohibited files in unencrypted or decrypted traffic	Protection
associated with malware detected in unencrypted or decrypted traffic	Malware

Table 36-2 License Requirements for Connection Logging in Access Control Policies (continued)

To log connections...	License
for traffic handled by user control or application control	Control
for traffic that the system filters using URL category and reputation data, and to display URL category and URL reputation information for URLs requested by monitored hosts	URL Filtering

Logging Security Intelligence (Blacklisting) Decisions

License: Protection

As a first line of defense against malicious Internet content, the ASA FirePOWER module includes the Security Intelligence feature, which allows you to immediately blacklist (block) connections based on the latest reputation intelligence, removing the need for a more resource-intensive, in-depth analysis. This traffic filtering takes place **before** any other policy-based inspection, analysis, or traffic handling.

Optionally, and recommended in passive deployments, you can use a monitor-only setting for Security Intelligence filtering. This allows the system to further analyze connections that would have been blacklisted, but still log the match to the blacklist.

Enabling Security Intelligence logging logs all blocked and monitored connections handled by an access control policy. However, the system does not log whitelist matches; logging of whitelisted connections depends on their eventual disposition.

When the system logs a connection event as the result of Security Intelligence filtering, it also logs a matching Security Intelligence event, which is a special kind of connection event that you can view and analyze separately. Both types of events use the **Action** and **Reason** fields to reflect the blacklist match. Additionally, so that you can identify the blacklisted IP address in the connection, host icons next to blacklisted and monitored IP addresses look slightly different in the event viewer.

Logging Blocked Blacklisted Connections

For a blocked connection, the system logs beginning-of-connection Security Intelligence and connection events. Because blacklisted traffic is immediately denied without further inspection, there is no unique end of connection to log. For these events, the action is `Block` and the reason is `IP Block`.

`IP Block` connection events have a threshold of 15 seconds per unique initiator-responder pair. That is, once the system generates an event when it blocks a connection, it does not generate another connection event for additional blocked connections between those two hosts for the next 15 seconds, regardless of port or protocol.

Logging Monitored Blacklisted Connections

For connections monitored—rather than blocked—by Security Intelligence, the system logs end-of-connection Security Intelligence and connection events to the ASA FirePOWER module. This logging occurs regardless of how the connection is later handled by an SSL policy, access control rule, or the access control default action.

For these connection events, the action depends on the connection's eventual disposition. The **Reason** field contains `IP Monitor`, as well as any other reason why the connection may have been logged.

Note that the system may also generate beginning-of-connection events for monitored connections, depending on the logging settings in the access control rule or default action that later handles the connection.

To log blacklisted connections:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to configure.

The access control policy editor appears.

Step 3 Select the Security Intelligence tab.

Security Intelligence settings for the access control policy appear.

Step 4 Click the logging icon ( ).

The Blacklist Options pop-up window appears.

Step 5 Select the **Log Connections** check box.

Step 6 Specify where to send connection and Security Intelligence events. You have the following choices:

- To send events to the ASA FirePOWER module, select **Event Viewer**.
- To send events to an external syslog server, select **Syslog**, then select a syslog alert response from the drop-down list. Optionally, you can add a syslog alert response by clicking the add icon ( ); see Creating a Syslog Alert Response, page 38-3.
- To send connection events to an SNMP trap server, select **SNMP Trap**, then select an SNMP alert response from the drop-down list. Optionally, you can add an SNMP alert response by clicking the add icon ( ); see Creating an SNMP Alert Response, page 38-2.

You **must** send events to the **Event Viewer** if you want to set blacklisted objects to monitor-only, or perform any other ASA FirePOWER module-based analysis on connection events generated by Security Intelligence filtering. For more information, see Logging Connections to the ASA FirePOWER Module or External Server, page 36-4.

Step 7 Click **OK** to set your logging options.

The Security Intelligence tab appears again.

Step 8 Click **Store ASA FirePOWER Changes**.

You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Logging Connections Based on Access Control Handling

License: Any

Within an access control policy, access control rules provide a granular method of handling network traffic. So that you can log only critical connections, you enable connection logging on a per-access-control-rule basis—if you enable connection logging for a rule, the system logs all connections handled by that rule.

You can also log connections for the traffic handled by the default action of your access control policy. The default action determines how the system handles traffic that matches none of the access control rules in the policy (except Monitor rules, which match and log—but do not handle or inspect—traffic).

Note that even if you disable logging for all access control rules and the default action, end-of-connection events may still be logged to the ASA FirePOWER module if the connection matches an access control rule and contains an intrusion attempt, prohibited file, or malware, or if it was decrypted by the system and you enabled logging for the connection in the SSL policy.

Depending on the rule or default policy action and the associated inspection options that you configure, your logging options differ. For more information, see:

- Logging Connections Matching an Access Control Rule, page 36-10
- Logging Connections Handled by the Access Control Default Action, page 36-11

Logging Connections Matching an Access Control Rule

License: Any

To log only critical connections, you enable connection logging on a per-access-control-rule basis. If you enable logging for a rule, the system logs all connections handled by that rule.

Depending on the rule action and intrusion and file inspection configuration of the rule, your logging options differ; see Understanding How Access Control and SSL Rule Actions Affect Logging, page 36-4. Also, note that even if you disable logging for an access control rule, end-of-connection events for connections matching that rule may still be logged to the ASA FirePOWER module if the connection:

- contains an intrusion attempt, prohibited file, or malware
- previously matched at least one access control Monitor rule

To configure an access control rule to log connection, file, and malware information:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to modify.

The access control policy editor appears, focused on the Rules tab.

Step 3 Click the edit icon ( ) next to the rule where you want to configure logging.

The access control rule editor appears.

Step 4 Select the Logging tab.

The Logging tab appears.

Step 5 Specify whether you want to **Log at Beginning and End of Connection**, **Log at End of Connection**, or you want **No Logging at Connection**.

For a single non-blocked connection, the end-of-connection event contains all of the information in the beginning-of-connection event, as well as information gathered over the duration of the session. Because blocked traffic is immediately denied without further inspection, the system logs only beginning-of-connection events for Block rules. For this reason, when you set the rule action to **Block** or **Block with reset** you are prompted **Log at Beginning of Connection**.

Step 6 Use the **Log Files** check box to specify whether the system should log any file and malware events associated with the connection.

The system automatically enables this option when you associate a file policy with the rule to perform either file control or AMP. Cisco recommends you leave this option enabled; see Disabling File and Malware Event Logging for Allowed Connections, page 36-7.

Step 7 Specify where to send connection events. You have the following choices:

- To send connection events to the ASA FirePOWER module, select **Event Viewer**. You cannot disable this option for Monitor rules.
- To send events to an external syslog server, select **Syslog**, then select a syslog alert response from the drop-down list. Optionally, you can add a syslog alert response by clicking the add icon ( ); see Creating a Syslog Alert Response, page 38-3.
- To send events to an SNMP trap server, select **SNMP Trap**, then select an SNMP alert response from the drop-down list. Optionally, you can add an SNMP alert response by clicking the add icon ( ); see Creating an SNMP Alert Response, page 38-2.

You **must** send events to the event viewer if you want to perform ASA FirePOWER module-based analysis on connection events. For more information, see Logging Connections to the ASA FirePOWER Module or External Server, page 36-4.

Step 8 Click **Store ASA FirePOWER Changes** to save the rule.

Your rule is saved. You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Logging Connections Handled by the Access Control Default Action

License: Any

You can log connections for the traffic handled by the default action of your access control policy. The default action determines how the system handles traffic that matches none of the access control rules in the policy (except Monitor rules, which match and log—but do not handle or inspect—traffic); see Setting Default Handling and Inspection for Network Traffic, page 4-4.

The mechanisms and options for logging connections handled by the policy default action largely parallel the options for logging connections handled by individual access control rules, as described in the following table. That is, except for blocked traffic, the system logs the beginning and end of connections, and you can send connection events to the ASA FirePOWER module, or to an external syslog or SNMP trap server.

Table 36-3 Access Control Default Action Logging Options

Default Action	Compare To	See...
Access Control: Block All Traffic	Block rules	Understanding Logging for Blocked and Interactively Blocked Connections, page 36-5
Access Control: Trust All Traffic	Trust rules	Understanding Logging for Trusted Connections, page 36-5
Intrusion Prevention	Allow rules with associated intrusion policies	Understanding Logging for Allowed Connections, page 36-6

However, there are some differences between logging connections handled by access control rules versus the default action:

- The default action has no file logging options. You cannot perform file control or AMP using the default action.
- When an intrusion policy associated with the access control default action generates an intrusion event, the system does **not** automatically log the end of the associated connection. This is useful for intrusion detection and prevention-only deployments, where you do not want to log any connection data.

  An exception to this rule occurs if you enable beginning- and end-of-connection logging for the default action. In that case, the system **does** log the end of the connection when an associated intrusion policy triggers, in addition to logging the beginning of the connection.

Note that even if you disable logging for the default action, end-of-connection events for connections matching that rule may still be logged to the ASA FirePOWER module if the connection previously matched at least one access control Monitor rule, or was inspected and logged by an SSL policy.

To log connections in traffic handled by the access control default action:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to modify.

The access control policy editor appears, focused on the Rules tab.

Step 3 Click the logging icon ( ) next to the **Default Action** drop-down list.

The Logging pop-up window appears.

Step 4 Specify whether you want to **Log at Beginning and End of Connection**, **Log at End of Connection**, or you want **No Logging at Connection**.

For a single non-blocked connection, the end-of-connection event contains all of the information in the beginning-of-connection event, as well as information gathered over the duration of the session. Because blocked traffic is immediately denied without further inspection, the system logs only beginning-of-connection events for the Block All Traffic default action. For this reason, when you set the default action to **Access Control: Block All Traffic** you are prompted **Log at Beginning of Connection**.

Step 5 Specify where to send connection events. You have the following choices:

- To send connection events to the ASA FirePOWER module, select **Event Viewer**. You cannot disable this option for Monitor rules.
- To send events to an external syslog server, select **Syslog**, then select a syslog alert response from the drop-down list. Optionally, you can add a syslog alert response by clicking the add icon ( ); see Creating a Syslog Alert Response, page 38-3.
- To send events to an SNMP trap server, select **SNMP Trap**, then select an SNMP alert response from the drop-down list. Optionally, you can add an SNMP alert response by clicking the add icon ( ); see Creating an SNMP Alert Response, page 38-2.

You **must** send events to the event viewer if you want to perform ASA FirePOWER module-based analysis on connection events. For more information, see Logging Connections to the ASA FirePOWER Module or External Server, page 36-4.

Step 6 Click **Store ASA FirePOWER Changes** to save the policy.

Your policy is saved. You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Logging URLs Detected in Connections

License: Any

When you log an end-of-connection event to the ASA FirePOWER module for HTTP traffic, the system records the URL requested by the monitored host during the session.

By default, the system stores the first 1024 characters of the URL in the connection log. However, you can configure the system to store up to 4096 characters per URL to make sure you capture the full URLs requested by monitored hosts. Or, if you are uninterested in the individual URLs visited, you can disable URL storage entirely by storing zero characters. Depending on your network traffic, disabling or limiting the number of stored URL characters may improve system performance.

Note that disabling URL logging does not affect URL filtering. Access control rules properly filter traffic based on requested URLs, their categories, and reputations, even though the system does not record the individual URLs requested in the traffic handled by those rules. For more information, see Blocking URLs, page 8-7.

To customize the number of URL characters you store:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to configure.

The access control policy editor appears.

Step 3 Select the Advanced tab.

Advanced settings for the access control policy appear.

Step 4 Click the edit icon ( ) next to General Settings.

The General Settings pop-up window appears.

Step 5 Type the **Maximum URL characters to store in connection events**.

You can specify any number from zero to 4096. Storing zero characters disables URL storage without disabling URL filtering.

Step 6 Click **OK**.

Advanced settings for the access control policy appear.

Step 7 Click **Store ASA FirePOWER Changes** to save the policy.

Your policy is saved. You must apply the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Logging Encrypted Connections

License: Any

As part of access control, the *SSL inspection* feature allows you to use an SSL policy to decrypt encrypted traffic for further evaluation by access control rules. You can force the system to log these decrypted connections, regardless of how the system later handles or inspects the traffic. You can also log connections when you block encrypted traffic, or when you allow it to pass to access control rules without decryption.

Connection logs for encrypted sessions contain details about the encryption, such as the certificate used to encrypt that session. You configure connection logging for encrypted sessions in the SSL policy on a per-SSL rule basis so that you only log critical connections.

For more information, see the following sections:

- Logging Decryptable Connections with SSL Rules, page 36-14
- Setting Default Logging for Encrypted and Undecryptable Connections, page 36-15

Logging Decryptable Connections with SSL Rules

License: Any

Within an SSL policy, *SSL rules* provide a granular method of handling encrypted traffic across multiple managed devices. So that you can log only critical connections, you enable connection logging on a per-SSL-rule basis—if you enable connection logging for a rule, the system logs all connections handled by that rule.

For encrypted connections inspected by an SSL policy, you can log connection events to an external syslog or SNMP trap server. You can log only end-of-connection events, however:

- for blocked connections (Block, Block with reset), the system immediately ends the session and generates an event
- for monitored connections (Monitor) and connections that you pass to access control rules (Decrypt, Do not decrypt), the system generates an event when the session ends, regardless of the logging configuration of the access control rule or default action that later handles it

For more information, see Understanding How Access Control and SSL Rule Actions Affect Logging, page 36-4.

To log decryptable connections:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > SSL**.

The SSL Policy page appears.

Step 2 Click the edit icon ( ) next to the rule where you want to configure logging.

The SSL rule editor appears.

Step 3 Select the Logging tab.

The Logging tab appears.

Step 4 Select **Log at End of Connection**.

Step 5 Specify where to send connection events. You have the following choices:

- To send events to an external syslog, select **Syslog**, then select a syslog alert response from the drop-down list. Optionally, you can add a syslog alert response by clicking the add icon ( ); see Creating a Syslog Alert Response, page 38-3.
- To send events to an SNMP trap server, select **SNMP Trap**, then select an SNMP alert response from the drop-down list. Optionally, you can add an SNMP alert response by clicking the add icon ( ); see Creating an SNMP Alert Response, page 38-2.

Step 6 Click **Add** to save your changes.

You must apply the access control policy the SSL policy is associated with for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Setting Default Logging for Encrypted and Undecryptable Connections

License: SSL

You can log connections for the traffic handled by the default action of your SSL policy. These logging settings also govern how the system logs undecryptable sessions.

The SSL policy default action determines how the system handles encrypted traffic that matches none of the SSL rules in the policy (except Monitor rules, which match and log—but do not handle or inspect—traffic). If your SSL policy does not contain any SSL rules, the default action determines how all encrypted sessions on your network are logged. For more information, see Setting Default Handling and Inspection for Encrypted Traffic, page 15-3.

You can configure the SSL policy default action to log connection events to an external syslog or SNMP trap server. You can log only end-of-connection events, however:

- for blocked connections (Block, Block with reset), the system immediately ends the sessions and generates an event
- for connections that you allow to pass unencrypted to access control rules (Do not decrypt), the system generates an event when the session ends

Note that even if you disable logging for the SSL policy default action, end-of-connection events may still be logged if the connection previously matched at least one SSL Monitor rule, or later matches an access control rule or the access control policy default action.

To set the default handling for encrypted and undecryptable traffic:

Access: Admin/Access Admin/Network Admin/Security Approver

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > SSL**.

The SSL Policy page appears.

Step 2 Click the logging icon ( ) next to the **Default Action** drop-down list.

The Logging pop-up window appears.

Step 3 Select **Log at End of Connection** to enable logging connection events.

Step 4 Specify where to send connection events. You have the following choices:

- To send events to an external syslog server, select **Syslog**, then select a syslog alert response from the drop-down list. Optionally, you can configure a syslog alert response by clicking the add icon ( ); see Creating a Syslog Alert Response, page 38-3.

- To send events to an SNMP trap server, select **SNMP Trap**, then select an SNMP alert response from the drop-down list. Optionally, you can configure an SNMP alert response by clicking the add icon ( ); see Creating an SNMP Alert Response, page 38-2.

Step 5 Click **OK** to save your changes.

You must apply the access control policy the SSL policy is associated with for your changes to take effect; see Deploying Configuration Changes, page 4-11.



Viewing Events

You can view real-time events logged against the traffic inspected by the ASA FirePOWER module.

Note The module only caches the most recent 100 events in memory.

For more information, see the following sections:

- Accessing ASA FirePOWER Real-Time Events, page 37-1
- Understanding ASA FirePOWER Event Types, page 37-2
- Event Fields in ASA FirePOWER Events, page 37-3
- Intrusion Rule Classifications, page 37-12

Accessing ASA FirePOWER Real-Time Events

You can view events detected by the ASA FirePOWER module in several predefined event views or create a custom event view to view the event fields you select.

Note The module only caches the most recent 100 events in memory.

To view ASA FirePOWER events:

Step 1 Select **Monitoring > ASA FirePOWER Monitoring > Real-time Eventing**.

Step 2 You have two choices:

- Click an existing tab for the type of event you want to view: connection events, security intelligence events, intrusion events, file events, or malware events.
- Click the + icon to create a custom event view and select the event fields you want to include in the view.

For more information, see Understanding ASA FirePOWER Event Types, page 37-2 and Event Fields in ASA FirePOWER Events, page 37-3.

Understanding ASA FirePOWER Event Types

The ASA FirePOWER module provides real-time event viewing of event fields from five event types: connection events, security intelligence events, intrusion events, file events, and malware events.

Connection Events

Connection logs, called *connection events*, contain data about the detected sessions. The information available for any individual connection event depends on several factors, but in general includes:

- basic connection properties: timestamp, source and destination IP address, ingress and egress zones, the device that handled the connection, and so on
- additional connection properties discovered or inferred by the system: applications, requested URLs, or users associated with the connection, and so on
- metadata about why the connection was logged: which access control rule (or other configuration) in which policy handled the traffic, whether the connection was allowed or blocked, and so on

Various settings in access control give you granular control over which connections you log, when you log them, and where you store the data. You can log any connection that your access control policies can successfully handle. You can enable connection logging in the following situations:

- when a connection is blacklisted (blocked) or monitored by the reputation-based Security Intelligence feature
- when a connection is handled by an access control rule or the access control default action

In addition to the logging that you configure, the system automatically logs most connections where the system detects a prohibited file, malware, or intrusion attempt.

Security Intelligence Events

When you enable Security Intelligence logging, blacklist matches automatically generate *Security Intelligence events* as well as connection events. A Security Intelligence event is a special kind of connection event that you can view and analyze separately. For detailed information on configuring connection logging, including Security Intelligence blacklisting decisions, see Logging Connections in Network Traffic, page 36-1.

Tip General information about connection events also pertains to Security Intelligence events, unless otherwise noted. For more information on Security Intelligence, see Blacklisting Using Security Intelligence IP Address Reputation, page 5-1.

Intrusion Events

The system examines the packets that traverse your network for malicious activity that could affect the availability, integrity, and confidentiality of a host and its data. When the system identifies a possible intrusion, it generates an *intrusion event*, which is a record of the date, time, type of exploit, and contextual information about the source of the attack and its target.

File Events

File events represent files that the system detected, and optionally blocked, in network traffic.

The system logs the file events generated when a managed device detects or blocks a file in network traffic, according to the rules in currently applied file policies.

Malware Events

Malware events represent malware files detected, and optionally blocked, in network traffic by the system.

With a Malware license, your ASA FirePOWER module can detect malware in network traffic as part of your overall access control configuration; see Understanding and Creating File Policies, page 35-4.

The following scenarios can lead to generating malware events:

- If a managed device detects one of a set of specific file types, the ASA FirePOWER module performs a malware cloud lookup, which returns a file disposition to the ASA FirePOWER module of `Malware`, `Clean`, or `Unknown`.
- If the ASA FirePOWER module cannot establish a connection with the cloud, or the cloud is otherwise unavailable, the file disposition is `Unavailable`. You may see a small percentage of events with this disposition; this is expected behavior.
- If the managed device detects a file on the clean list, the ASA FirePOWER module assigns a file disposition of `Clean` to the file.

The ASA FirePOWER module logs records of files' detection and dispositions, along with other contextual data, as malware events.

Files detected in network traffic and identified as malware by the ASA FirePOWER module generate both a file event and a malware event. This is because to detect malware in a file, the system must first detect the file itself.

Event Fields in ASA FirePOWER Events

Action

For connection or security intelligence events, the action associated with the access control rule or default action that logged the connection:

- `Allow` represents explicitly allowed and user-bypassed interactively blocked connections.
- `Trust` represents trusted connections. TCP connections detected by a trust rule on the first packet only generate an end-of-connection event. The system generates the event one hour after the final session packet.
- `Block` and `Block with reset` represent blocked connections. The system also associates the `Block` action with connections blacklisted by Security Intelligence, connections where an exploit was detected by an intrusion policy, and connections where a file was blocked by a file policy.
- `Interactive Block` and `Interactive Block with reset` mark the beginning-of-connection event that you can log when the system initially blocks a user's HTTP request using an Interactive Block rule. If the user clicks through the warning page that the system displays, any additional connection events you log for the session have an action of `Allow`.
- `Default Action` indicates the connection was handled by the default action.
- For Security Intelligence-monitored connections, the action is that of the first non-Monitor access control rule triggered by the connection, or the default action. Similarly, because traffic matching a Monitor rule is always handled by a subsequent rule or by the default action, the action associated with a connection logged due to a monitor rule is never `Monitor`.

For file or malware events, the file rule action associated with the rule action for the rule the file matched, and any associated file rule action options.

Allowed Connection

Whether the system allowed the traffic flow for the event.

Application

The application detected in the connection.

Application Business Relevance

The business relevance associated with the application traffic detected in the connection: `Very High`, `High`, `Medium`, `Low`, or `Very Low`. Each type of application detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Application Categories

Categories that characterize the application to help you understand the application's function.

Application Risk

The risk associated with the application traffic detected in the connection: `Very High`, `High`, `Medium`, `Low`, or `Very Low`. Each type of application detected in the connection has an associated risk; this field displays the highest of those.

Application Tag

Tags that characterize the application to help you understand the application's function.

Block Type

The type of block specified in the access control rule matching the traffic flow in the event: block or interactive block.

Client

The client application detected in the connection.

If the system cannot identify the specific client used in the connection, this field displays `client` appended to the application protocol name to provide a generic name, for example, `FTP client`.

Client Business Relevance

The business relevance associated with the client traffic detected in the connection: `Very High`, `High`, `Medium`, `Low`, or `Very Low`. Each type of client detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Client Categories

Categories that characterize the client detected in the traffic to help you understand the client’s function.

Client Risk

The risk associated with the client traffic detected in the connection: `Very High`, `High`, `Medium`, `Low`, or `Very Low`. Each type of client detected in the connection has an associated risk; this field displays the highest of those.

Client Tag

Tags that characterize the client detected in the traffic to help you understand the client’s function.

Client Version

The version of the client detected in the connection.

Connection

The unique ID for the traffic flow, internally generated.

Connection Blocktype Indicator

The type of block specified in the access control rule matching the traffic flow in the event: block or interactive block.

Connection Bytes

The total bytes for the connection.

Connection Time

The time for the beginning of the connection.

Connection Timestamp

The time the connection was detected.

Context

The metadata identifying the security context through which the traffic passed. Note that the system only populates this field for devices in multiple context mode.

Denied Connection

Whether the system denied the traffic flow for the event.

Destination Country and Continent

The country and continent of the receiving host.

Destination IP

The IP address used by the receiving host.

Destination Port, Destination Port Icode, Destination Port/ICMP Code

The destination port or ICMP code used by the session responder.

Direction

The direction of transmission for a file.

Disposition

One of the following file dispositions:

- `Malware` indicates that the cloud categorized the file as malware.
- `Clean` indicates that the cloud categorized the file as clean, or that a user added the file to the clean list.
- `Unknown` indicates that a malware cloud lookup occurred before the cloud assigned a disposition. The file is uncategorized.
- `Custom Detection` indicates that a user added the file to the custom detection list.

- `Unavailable` indicates that the ASA FirePOWER module could not perform a malware cloud lookup. You may see a small percentage of events with this disposition; this is expected behavior.
- `N/A` indicates a Detect Files or Block Files rule handled the file and the ASA FirePOWER module did not perform a malware cloud lookup.

Egress Interface

The egress interface associated with the connection. Note that, if your deployment includes an asynchronous routing configuration, the ingress and egress interface may belong to the same interface set.

Egress Security Zone

The egress security zone associated with the connection.

Event

The event type.

Event Microseconds

The time, in microseconds, when the event was detected.

Event Seconds

The time, in seconds, when the event was detected.

Event Type

The type of event.

File Category

The general categories of file type, for example: `Office Documents`, `Archive`, `Multimedia`, `Executables`, `PDF files`, `Encoded`, `Graphics`, or `System Files`.

File Event Timestamp

The time and date the file or malware file was created.

File Name

The name of the file or malware file.

File SHA256

The SHA-256 hash value of the file.

File Size

The size of the file or malware file, in kilobytes.

File Type

The file type of the file or malware file, for example, `HTML` or `MSEXE`.

File/Malware Policy

The file policy associated with the generation of the event.

Filelog Blocktype Indicator

The type of block specified in the file rule matching the traffic flow in the event: block or interactive block.

Firewall Policy Rules/SI Category

The name of the blacklisted object that represents or contains the blacklisted IP address in the connection. The Security Intelligence category can be the name of a network object or group, the global blacklist, a custom Security Intelligence list or feed, or one of the categories in the Intelligence Feed. Note that this field is only populated if the **Reason** is `IP Block` or `IP Monitor`; entries in Security Intelligence event views always display a reason.

Firewall Rule

The access control rule or default action that handled the connection, as well as up to eight Monitor rules matched by that connection.

First Packet

The date and time the first packet of the session was seen.

HTTP Referrer

The HTTP referrer, which represents the referrer of a requested URL for HTTP traffic detected in the connection (such as a website that provided a link to, or imported a link from, another URL).

IDS Classification

The classification where the rule that generated the event belongs. See the Rule Classifications table for a list of rule classification names and numbers.

Impact

The impact level in this field indicates the correlation between intrusion data, network discovery data, and vulnerability information.

Impact Flag

See Impact.

Ingress Interface

The ingress interface associated with the connection. Note that, if your deployment includes an asynchronous routing configuration, the ingress and egress interface may belong to the same interface set.

Ingress Security Zone

The ingress security zone associated with the connection.

Initiator Bytes

The total number of bytes transmitted by the session initiator.

Initiator Country and Continent

When a routable IP is detected, the country and continent associated with the host IP address that initiated the session.

Initiator IP

The host IP address (and host name, if DNS resolution is enabled) that initiated the session responder.

Initiator Packets

The total number of packets transmitted by the session initiator.

Inline Result

One of the following:

- a black down arrow, indicating that the system dropped the packet that triggered the rule
- a gray down arrow, indicating that IPS would have dropped the packet if you enabled the **Drop when Inline** intrusion policy option (in an inline deployment), or if a Drop and Generate rule generated the event while the system was pruning
- blank, indicating that the triggered rule was not set to Drop and Generate Events
- Note that the system does not drop packets in a passive deployment, including when an inline interface is in tap mode, regardless of the rule state or the inline drop behavior of the intrusion policy.

IPS Blocktype Indicator

The action of the intrusion rule matching the traffic flow in the event.

Last Packet

The date and time the last packet of the session was seen.

MPLS Label

The Multiprotocol Label Switching label associated with the packet that triggered this intrusion event.

Malware Blocktype Indicator

The type of block specified in the file rule matching the traffic flow in the event: block or interactive block.

Message

The explanatory text for the event.

For rule-based intrusion events, the event message is pulled from the rule. For decoder- and preprocessor-based events, the event message is hard coded.

For malware events, any additional information associated with the malware event. For network-based malware events, this field is populated only for files whose disposition has changed.

Monitor Rules

Up to eight Monitor rules matched by that connection.

Netbios Domain

The NetBIOS domain used in the session.

Original Client Country and Continent

The country where the original client IP address belongs. To obtain this value, the system extracts the original client IP address from an X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header, then maps it to the country using the geolocation database (GeoDB). To populate this field, you must enable an access control rule that handles proxied traffic based on its original client.

Original Client IP

The original client IP address from an X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header. To populate this field, you must enable an access control rule that handles proxied traffic based on its original client.

Policy

The access control, intrusion, or network analysis policy (NAP), if any, associated with the generation of the event.

Policy Revision

The revision of the access control, file, intrusion, or network analysis policy (NAP), if any, associated with the generation of the event.

Priority

The event priority as determined by the Cisco VRT.

Protocol

The protocol detected in the connection.

Reason

The reason or reasons the connection was logged, in the following situations:

- `User Bypass` indicates that the system initially blocked a user's HTTP request, but the user chose to continue to the originally requested site by clicking through a warning page. A reason of `User Bypass` is always paired with an action of `Allow`.
- `IP Block` indicates that the system denied the connection without inspection, based on Security Intelligence data. A reason of `IP Block` is always paired with an action of `Block`.
- `IP Monitor` indicates that the system would have denied the connection based on Security Intelligence data, but you configured the system to monitor, rather than deny, the connection.
- `File Monitor` indicates that the system detected a particular type of file in the connection.
- `File Block` indicates the connection contained a file or malware file that the system prevented from being transmitted. A reason of `File Block` is always paired with an action of `Block`.
- `File Custom Detection` indicates the connection contained a file on the custom detection list that the system prevented from being transmitted.
- `File Resume Allow` indicates that file transmission was originally blocked by a Block Files or Block Malware file rule. After a new access control policy was applied that allowed the file, the HTTP session automatically resumed. Note that this reason only appears in inline deployments.
- `File Resume Block` indicates that file transmission was originally allowed by a Detect Files or Malware Cloud Lookup file rule. After a new access control policy was applied that blocked the file, the HTTP session automatically stopped. Note that this reason only appears in inline deployments.

- `Intrusion Block` indicates the system blocked or would have blocked an exploit (intrusion policy violation) detected in the connection. A reason of `Intrusion Block` is paired with an action of `Block` for blocked exploits and `Allow` for would-have-blocked exploits.
- `Intrusion Monitor` indicates the system detected, but did not block, an exploit detected in the connection. This occurs when the state of the triggered intrusion rule is set to **Generate Events**.
- `Content Restriction` indicates the system modified the packet to enforce content restrictions related to either the Safe Search or YouTube EDU feature.

Receive Times

The time the destination host or responder responded to the event.

Referenced Host

If the protocol in the connection is DNS, HTTP, or HTTPS, this field displays the host name that the respective protocol was using.

Responder Bytes

The total number of bytes transmitted by the session responder.

Responder Country and Continent

When a routable IP is detected, the country and continent associated with the host IP address for the session responder.

Responder Packets

The total number of packets transmitted by the session responder.

Responder IP

The host IP address (and host name, if DNS resolution is enabled) that responded to the session initiator.

Security Group Tag Name

The Security Group Tag (SGT) attribute of the packet involved in the connection. The SGT specifies the privileges of a traffic source within a trusted network. Security Group Access (a feature of both Cisco TrustSec and Cisco ISE) applies the attribute as packets enter the network.

Signature

The signature ID of the intrusion rule matching the traffic for the event.

Source Country and Continent

The country and continent of the sending host.

Source IP

The IP address used by the sending host in an intrusion event.

Source or Destination

The host originating or receiving the connection for the event.

Source Port, Source Port Type, Source Port/ICMP Type

The source port or ICMP type used by the session initiator.

TCP Flags

The TCP flags detected in the connection.

URL

The URL requested by the monitored host during the session.

URL Category

The category associated with the URL requested by the monitored host during the session, if available.

URL Reputation

The reputation associated with the URL requested by the monitored host during the session, if available.

URL Reputation Score

The reputation score associated with the URL requested by the monitored host during the session, if available.

User

The user of the host (**Receiving IP**) where the event occurred.

User Agent

User agent application information extracted from HTTP traffic detected in the connection.

VLAN

The innermost VLAN ID associated with the packet that triggered the event.

Web App Business Relevance

The business relevance associated with the web application traffic detected in the connection: `Very High`, `High`, `Medium`, `Low`, or `Very Low`. Each type of web application detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Web App Categories

Categories that characterize the web application detected in the traffic to help you understand the web application's function.

Web App Risk

The risk associated with the web application traffic detected in the connection: `Very High`, `High`, `Medium`, `Low`, or `Very Low`. Each type of web application detected in the connection has an associated risk; this field displays the highest of those.

Web App Tag

Tags that characterize the web application detected in the traffic to help you understand the web application's function.

Web Application

The web application detected in the traffic.

Intrusion Rule Classifications

Intrusion rules include an attack classification. The following table lists the name and number for each classification

Table 37-1 Rule Classifications

Number	Classification Name	Description
1	not-suspicious	Not Suspicious Traffic
2	unknown	Unknown Traffic
3	bad-unknown	Potentially Bad Traffic
4	attempted-recon	Attempted Information Leak
5	successful-recon-limited	Information Leak
6	successful-recon-largescale	Large Scale Information Leak
7	attempted-dos	Attempted Denial of Service
8	successful-dos	Denial of Service
9	attempted-user	Attempted User Privilege Gain
10	unsuccessful-user	Unsuccessful User Privilege Gain
11	successful-user	Successful User Privilege Gain
12	attempted-admin	Attempted Administrator Privilege Gain
13	successful-admin	Successful Administrator Privilege Gain
14	rpc-portmap-decode	Decode of an RPC Query
15	shellcode-detect	Executable Code was Detected
16	string-detect	A Suspicious String was Detected
17	suspicious-filename-detect	A Suspicious Filename was Detected
18	suspicious-login	An Attempted Login Using a Suspicious Username was Detected
19	system-call-detect	A System Call was Detected
20	tcp-connection	A TCP Connection was Detected
21	trojan-activity	A Network Trojan was Detected
22	unusual-client-port-connection	A Client was Using an Unusual Port
23	network-scan	Detection of a Network Scan
24	denial-of-service	Detection of a Denial of Service Attack
25	non-standard-protocol	Detection of a Non-Standard Protocol or Event
26	protocol-command-decode	Generic Protocol Command Decode
27	web-application-activity	Access to a Potentially Vulnerable Web Application
28	web-application-attack	Web Application Attack
29	misc-activity	Misc Activity
30	misc-attack	Misc Attack
31	icmp-event	Generic ICMP Event

Table 37-1 Rule Classifications

Number	Classification Name	Description
32	inappropriate-content	Inappropriate Content was Detected
33	policy-violation	Potential Corporate Privacy Violation
34	default-login-attempt	Attempt to Login By a Default Username and Password
35	sdf	Sensitive Data
36	malware-cnc	Known malware command and control traffic
37	client-side-exploit	Known client side exploit attempt
38	file-format	Known malicious file or file based exploit



Configuring External Alerting

While the ASA FirePOWER module provides various views of events within the module interface, you may want to configure external event notification to facilitate constant monitoring of critical systems. You can configure the module to generate alerts that notify you via email, SNMP trap, or syslog when one of the following is generated:

- a network-based malware event or retrospective malware event
- a connection event, triggered by a specific access control rule

To have the ASA FirePOWER module send these alerts, you must first create an *alert response*, which is a set of configurations that allows the module to interact with the external system where you plan to send the alert. Those configurations may specify, for example, SNMP alerting parameters or syslog facilities and priorities.

After you create the alert response, you associate it with the event that you want to use to trigger the alert. Note that the process for associating alert responses with events is different depending on the type of event:

- You associate alert responses with malware events using their own configuration pages.
- You associate SNMP and syslog alert responses with logged connections using access control rules and policies.

There is another type of alerting you can perform in the ASA FirePOWER module, which is to configure SNMP and syslog intrusion event notifications for individual intrusion events. You configure these notifications in intrusion policies; see Configuring External Alerting for Intrusion Rules, page 39-1 and Adding SNMP Alerts, page 27-31. The following table explains the licenses you must have to generate alerts.

Table 38-1 License Requirements for Generating Alerts

To generate an alert based on...	You need this license...
an intrusion event	Protection
a network-based malware event	Malware
a connection event	the license required to log the connection

For more information, see:

- Working with Alert Responses, page 38-2
- Logging Connections in Network Traffic, page 36-1

Working with Alert Responses

License: Any

The first step in configuring external alerting is to create an alert response, which is a set of configurations that allows the ASA FirePOWER module to interact with the external system where you plan to send the alert. You can create alert responses to send alerts via email, a simple network management protocol (SNMP) trap, or a system log (syslog).

The information you receive in an alert depends on the type of event that triggered the alert.

When you create an alert response, it is automatically enabled. Only enabled alert responses can generate alerts. To stop alerts from being generated, you can temporarily disable alert responses rather than deleting your configurations.

You manage alert responses on the Alerts page (**ASA FirePOWER Configuration > Policies > Actions Alerts**). The slider next to each alert response indicates whether it is active; only enabled alert responses can generate alerts. The page also indicates whether the alert response is being used in a configuration, for example, to log connections in an access control rule. You can sort alert responses by name, type, in use status, and enabled/disabled status by clicking the appropriate column header; click the column header again to reverse the sort.

For more information, see:

- Creating an SNMP Alert Response, page 38-2
- Creating a Syslog Alert Response, page 38-3
- Modifying an Alert Response, page 38-5
- Deleting an Alert Response, page 38-6
- Enabling and Disabling Alert Responses, page 38-6

Creating an SNMP Alert Response

License: Any

You can create SNMP alert responses using SNMPv1, SNMPv2, or SNMPv3.

Note If you want to monitor 64-bit values with SNMP, you must use SNMPv2 or SNMPv3. SNMPv1 does not support 64-bit monitoring.

To create an SNMP alert response:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Actions Alerts**.

The Alerts page appears.

Step 2 From the **Create Alert** drop-down menu, select **Create SNMP Alert**.

The Create SNMP Alert Configuration pop-up window appears.

Step 3 In the **Name** field, type the name that you want to use to identify the SNMP response.

Step 4 In the **Trap Server** field, type the hostname or IP address of the SNMP trap server, using alphanumeric characters.

Note that the system does **not** warn you if you enter an invalid IPv4 address (such as 192.169.1.456) in this field. Instead, the invalid address is treated as a hostname.

Step 5 From the **Version** drop-down list, select the SNMP version you want to use.

SNMP v3 is the default. If you select SNMP v1 or SNMP v2, different options appear.

Step 6 Which version of SNMP did you select?

- For SNMP v1 or SNMP v2, type the SNMP community name, using alphanumeric characters or the special characters `*` or `$`, in the **Community String** field and skip to step 12.
- For SNMP v3, type the name of the user that you want to authenticate with the SNMP server in the **User Name** field and continue with the next step.

Step 7 From the **Authentication Protocol** drop-down list, select the protocol you want to use for authentication.

Step 8 In the **Authentication Password** field, type the password required for authentication with the SNMP server.

Step 9 From the **Privacy Protocol** list, select **None** to use no privacy protocol or **DES** to use Data Encryption Standard as the privacy protocol.

Step 10 In the **Privacy Password** field, type the privacy password required by the SNMP server.

Step 11 In the **Engine ID** field, type an identifier for the SNMP engine, in hexadecimal notation, using an even number of digits.

When you use SNMPv3, the system uses an Engine ID value to encode the message. Your SNMP server requires this value to decode the message.

Cisco recommends that you use the hexadecimal version of the ASA FirePOWER module's IP address. For example, if the ASA FirePOWER module has an IP address of `10.1.1.77`, use `0a01014D0`.

Step 12 Click **Store ASA FirePOWER Changes**.

The alert response is saved and is automatically enabled.

Creating a Syslog Alert Response

License: Any

When configuring a syslog alert response, you can specify the severity and facility associated with the syslog messages to ensure that they are processed properly by the syslog server. The facility indicates the subsystem that creates the message and the severity defines the severity of the message. Facilities and severities are not displayed in the actual message that appears in the syslog, but are instead used to tell the system that receives the syslog message how to categorize it.



Tip For more detailed information about how syslog works and how to configure it, refer to the documentation for your system. On UNIX systems, the `man` pages for `syslog` and `syslog.conf` provide conceptual information and configuration instructions.

Although you can select any type of facility when creating a syslog alert response, you should select one that makes sense based on your syslog server; not all syslog servers support all facilities. For UNIX syslog servers, the `syslog.conf` file should indicate which facilities are saved to which log files on the server.

The following table lists the syslog facilities you can select.

Table 38-2 Available Syslog Facilities

Facility	Description
ALERT	An alert message.
AUDIT	A message generated by the audit subsystem.
AUTH	A message associated with security and authorization.
AUTHPRIV	A restricted access message associated with security and authorization. On many systems, these messages are forwarded to a secure file.
CLOCK	A message generated by the clock daemon. Note that syslog servers running a Windows operating system will use the `CLOCK` facility.
CRON	A message generated by the clock daemon. Note that syslog servers running a Linux operating system will use the `CRON` facility.
DAEMON	A message generated by a system daemon.
FTP	A message generated by the FTP daemon.
KERN	A message generated by the kernel. On many systems, these messages are printed to the console when they appear.
LOCAL0-LOCAL7	A message generated by an internal process.
LPR	A message generated by the printing subsystem.
MAIL	A message generated by a mail system.
NEWS	A message generated by the network news subsystem.
NTP	A message generated by the NTP daemon.
SYSLOG	A message generated by the syslog daemon.
USER	A message generated by a user-level process.
UUCP	A message generated by the UUCP subsystem.

The following table lists the standard syslog severity levels you can select.

Table 38-3 Syslog Severity Levels

Level	Description
ALERT	A condition that should be corrected immediately.
CRIT	A critical condition.
DEBUG	Messages that contain debugging information.
EMERG	A panic condition broadcast to all users.
ERR	An error condition.
INFO	Informational messages.
NOTICE	Conditions that are not error conditions, but require attention.
WARNING	Warning messages.

Before you start sending syslog alerts, make sure that the syslog server can accept remote messages.

To create a syslog alert:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Actions Alerts**.

The Alerts page appears.From the **Create Alert** drop-down menu, select **Create Syslog Alert**.

The Create Syslog Alert Configuration pop-up window appears.

Step 2 In the **Name** field, type the name you want to use to identify the saved response.

Step 3 In the **Host** field, type the hostname or IP address of your syslog server.

Note that the system does **not** warn you if you enter an invalid IPv4 address (such as 192.168.1.456) in this field. Instead, the invalid address is treated as a hostname.

Step 4 In the **Port** field, type the port the server uses for syslog messages.

By default, this value is 514.

Step 5 From the **Facility** list, select a facility.

See the Available Syslog Facilities table for a list of the available facilities.

Step 6 From the **Severity** list, select a severity.

See the Syslog Severity Levels table for a list of the available severities.

Step 7 In the **Tag** field, type the tag name that you want to appear with the syslog message.

Use only alphanumeric characters in tag names. You **cannot** use spaces or underscores.

As an example, if you wanted all messages sent to the syslog to be preceded with `FromMC`, type `FromMC` in the field.

Step 8 Click **Store ASA FirePOWER Changes**.

The alert response is saved and is automatically enabled.

Modifying an Alert Response

License: Any

For most types of alerting, if an alert response is enabled and in use, changes to the alert response take effect immediately. However, for alert responses used in access control rules to log connection events, changes do not take effect until you reapply the access control policy.

To edit an alert response:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Actions Alerts**.

The Alerts page appears.

Step 2 Next to the alert response you want to edit, click the edit icon ( ).

A configuration pop-up window for that alert response appears.

Step 3 Make changes as needed.

Step 4 Click **Store ASA FirePOWER Changes**.

The alert response is saved.

Deleting an Alert Response

License: Any

You can delete any alert response that is not in use.

To delete an alert response:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Actions Alerts**.

The Alerts page appears.

Step 2 Next to the alert response you want to delete, click the delete icon ( ).

Step 3 Confirm that you want to delete the alert response.

The alert response is deleted.

Enabling and Disabling Alert Responses

License: Any

Only enabled alert responses can generate alerts. To stop alerts from being generated, you can temporarily disable alert responses rather than deleting your configurations. Note that if an alert is in use when you disable it, it is still considered in use even though it is disabled.

To enable or disable an alert response:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Actions Alerts**.

The Alerts page appears.

Step 2 Next to the alert response you want to enable or disable, click the enable/disable slider.

If the alert response was enabled, it is disabled. If it was disabled, it is enabled.



Configuring External Alerting for Intrusion Rules

While the ASA FirePOWER module provides various views of intrusion events within the user interface, some enterprises prefer to define external intrusion event notification to facilitate constant monitoring of critical systems. You can enable logging to syslog facilities or send event data to an SNMP trap server.

Within each intrusion policy, you can specify intrusion event notification limits, set up intrusion event notification to external logging facilities, and configure external responses to intrusion events.

Tip Some analysts prefer not to receive multiple alerts for the same intrusion event, but want to control how often they are notified of a given intrusion event occurrence. See Filtering Intrusion Event Notification Per Policy, page 27-21 for more information.

There is another type of alerting you can perform in the ASA FirePOWER module, outside of your intrusion policies. You can configure SNMP and syslog alert responses for other types of events, including connection events logged by specific access control rules. For more information, see Configuring External Alerting, page 38-1.

See the following sections for more information on external intrusion event notification:

- Using SNMP Responses, page 39-1 describes the options you can configure to send event data to specified SNMP trap servers and provides the procedure for specifying the SNMP alerting options.
- Using Syslog Responses, page 39-4 describes the options you can configure to send event data to an external syslog and provides the procedure for specifying the syslog alerting options.

Using SNMP Responses

License: Protection

An *SNMP trap* is a network management notification. You can configure the device to send intrusion event notifications as SNMP traps, also known as *SNMP alerts*. Each SNMP alert includes:

- the name of the server generating the trap
- the IP address of the device that detected it
- the name of the device that detected it
- the event data

You can set a variety of SNMP alerting parameters. Available parameters vary depending on the version of SNMP you use. For details on enabling and disabling SNMP alerting, see Configuring Advanced Settings in an Intrusion Policy, page 26-6.

Tip If your network management system requires a management information base file (MIB), you can obtain it from the ASA FirePOWER module at `/etc/sf/DCEALERT.MIB`.

SNMP v2 Options

For SNMP v2, you can specify the options described in the following table.

Table 39-1 SNMP v2 Options

Option	Description
Trap Type	The trap type to use for IP addresses that appear in the alerts. If your network management system correctly renders the INET_IPV4 address type, then you can select **as Binary**. Otherwise, select **as String**. For example, HP Openview requires the string type.
Trap Server	The server that will receive SNMP traps notification. You can specify a single IP address or hostname.
Community String	The community name.

SNMP v3 Options

For SNMP v3, you can specify the options described in the following table.

Note When using SNMP v3, the appliance uses an Engine ID value to encode the message. Your SNMP server requires this value to decode the message. Currently, this Engine ID value will always be the hexadecimal version of the appliance's IP address with `01` at the end of the string. For example, if the appliance sending the SNMP alert has an IP address of `172.16.1.50`, the Engine ID is `0xAC10013201` or, if the appliance has an IP address of `10.1.1.77`, `0x0a01014D01` is used as the Engine ID.

Table 39-2 SNMP v3 Options

Option	Description
Trap Type	The trap type to use for IP addresses that appear in the alerts. If your network management system correctly renders the INET_IPV4 address type, then you can select **as Binary**. Otherwise, select **as String**. For example, HP Openview requires the string type.
Trap Server	The server that will receive SNMP traps notification. You can specify a single IP address or hostname.
Authentication Password	The password required for authentication. SNMP v3 uses either the Message Digest 5 (MD5) hash function or the Secure Hash Algorithm (SHA) hash function to encrypt this password, depending on configuration. If you specify an authentication password, authentication is enabled.

Table 39-2 SNMP v3 Options (continued)

Option	Description
Private Password	The SNMP key for privacy. SNMP v3 uses the Data Encryption Standard (DES) block cipher to encrypt this password. If you specify a private password, privacy is enabled. If you specify a private password, you must also specify an authentication password.
User Name	Your SNMP user name.

For information about configuring SNMP Alerting, see Configuring SNMP Responses, page 39-3.

Configuring SNMP Responses

License: Protection

You can configure SNMP alerting in an intrusion policy. After you apply the policy as part of an access control policy, the system notifies you of any intrusion events it detects via SNMP trap. For more details on SNMP alerting, see Using SNMP Responses, page 39-1.

To configure SNMP alerting options:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy**.

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

Step 4 You have two choices, depending on whether **SNMP Alerting** under External Responses is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The SNMP Alerting page appears.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 5 Specify the trap type format that you want to use for IP addresses that appear in the alerts, **as Binary** or **as String**.

Note If your network management system correctly renders the INET_IPV4 address type, then you can use the **as Binary** option. Otherwise, use the **as String** option. For example, HP OpenView requires the **as String** option.

Step 6 Select either SNMP v2 or SNMP v3:

- To configure SNMP v2, enter the IP address and the community name of the trap server you want to use in the corresponding fields. See SNMP v2 Options, page 39-2.
- To configure SNMP v3, enter the IP address of the trap server you want to use, an authentication password, a private password, and a user name in the corresponding fields. See SNMP v3 Options, page 39-2 for more information.

Note You must select SNMP v2 **or** SNMP v3.

Note When you enter an SNMP v3 password, the password displays in plain text during initial configuration but is saved in encrypted format.

Step 7 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.

Using Syslog Responses

License: Protection

The system log, or *syslog,* is the standard logging mechanism for network event logging. You can send *syslog alerts*, which are intrusion event notifications, to the syslog on an appliance. The syslog allows you to categorize information in the syslog by priority and facility. The *priority* reflects the severity of the alert and the *facility* indicates the subsystem that generated the alert. Facilities and priorities are not displayed in the actual message that appears in syslog, but are instead used to tell the system that receives the syslog message how to categorize it.

Syslog alerts contain the following information:

- date and time of alert generation
- event message
- event data
- generator ID of the triggering event
- Snort ID of the triggering event
- revision

In an intrusion policy, you can turn on syslog alerting and specify the syslog priority and facility associated with intrusion event notifications in the syslog. When you apply the intrusion policy as part of an access control policy, the system then sends syslog alerts for the intrusion events it detects to the syslog facility on the local host or on the logging host specified in the policy. The host receiving the alerts uses the facility and priority information you set when configuring syslog alerting to categorize the alerts.

The following table lists the facilities you can select when configuring syslog alerting. Be sure to configure a facility that makes sense based on the configuration of the remote syslog server you use. The `syslog.conf` file located on the remote system (if you are logging syslog messages to a UNIX- or Linux-based system) indicates which facilities are saved to which log files on the server.

Table 39-3 Available Syslog Facilities

Facility	Description
AUTH	A message associated with security and authorization.
AUTHPRIV	A restricted access message associated with security and authorization. On many systems, these messages are forwarded to a secure file.
CRON	A message generated by the clock daemon.
DAEMON	A message generated by a system daemon.
FTP	A message generated by the FTP daemon.
KERN	A message generated by the kernel. On many systems, these messages are printed to the console when they appear.
LOCAL0-LOCAL7	A message generated by an internal process.
LPR	A message generated by the printing subsystem.
MAIL	A message generated by a mail system.
NEWS	A message generated by the network news subsystem.
SYSLOG	A message generated by the syslog daemon.
USER	A message generated by a user-level process.
UUCP	A message generated by the UUCP subsystem.

Select one of the following standard syslog priority levels to display on all notifications generated by this alert:

Table 39-4 Syslog Priority Levels

Level	Description
EMERG	A panic condition broadcast to all users
ALERT	A condition that should be corrected immediately
CRIT	A critical condition
ERR	An error condition
WARNING	Warning messages
NOTICE	Conditions that are not error conditions, but require attention
INFO	Informational messages
DEBUG	Messages that contain debug information

For more detailed information about how syslog works and how to configure it, refer to the documentation that accompanies your system. If you are logging to a UNIX- or Linux-based system's syslog, the `syslog.conf` man file (type `man syslog.conf` at the command line) and syslog man file (type `man syslog` at the command line) provide information about how syslog works and how to configure it.

Configuring Syslog Responses

License: Protection

You can configure syslog alerting in an intrusion policy. After you apply the policy as part of an access control policy, the system notifies you of any intrusion events it detects via the syslog. For more information on syslog alerting, see Using Syslog Responses, page 39-4.

To configure syslog alerting options:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies> Intrusion Policy**.

The Intrusion Policy page appears.

Step 2 Click the edit icon ( ) next to the policy you want to edit.

If you have unsaved changes in another policy, click **OK** to discard those changes and continue. See Resolving Conflicts and Committing Policy Changes, page 18-15 for information on saving unsaved changes in another policy.

The Policy Information page appears.

Step 3 Click **Advanced Settings** in the navigation panel on the left.

The Advanced Settings page appears.

Step 4 You have two choices, depending on whether **Syslog Alerting** under External Responses is enabled:

- If the configuration is enabled, click **Edit**.
- If the configuration is disabled, click **Enabled**, then click **Edit**.

The Syslog Alerting page appears.

A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. See Using Layers in a Network Analysis or Intrusion Policy, page 19-1 for more information.

Step 5 Optionally, in the **Logging Hosts** field, enter the remote access IP address you want to specify as logging host. Separate multiple hosts with commas.

Step 6 Select facility and priority levels from the drop-down lists.

See Using Syslog Responses, page 39-4 for details on facility and priority options.

Step 7 Save your policy, continue editing, discard your changes, revert to the default configuration settings in the base policy, or exit while leaving your changes in the system cache. See Resolving Conflicts and Committing Policy Changes, page 18-15 for more information.



Using the ASA FirePOWER Dashboard

The ASA FirePOWER module dashboard provides you with at-a-glance views of current system status. The dashboard displays widgets in a three-column layout. Widgets are small, self-contained components that provide insight into different aspects of the ASA FirePOWER module. Your system is delivered with several predefined widgets. For example, the Appliance Information widget tells you the appliance name, model, and currently running version of the ASA FirePOWER module software.

The dashboard has a time range that constrains its widgets. You can change the time range to reflect a period as short as the last hour or as long as the last year.

Each appliance is delivered with a default dashboard. This dashboard provides the user with general system status information for your ASA FirePOWER module deployment.

For more information on the dashboard and its contents, see the following sections:

- Understanding Dashboard Widgets, page 40-1
- Understanding the Predefined Widgets, page 40-2
- Working with the Dashboard, page 40-5

Understanding Dashboard Widgets

License: Any

The dashboard displays multiple widgets in a three-column layout. The ASA FirePOWER module is delivered with several predefined dashboard widgets, each of which provides insight into a different aspect of the system. You can minimize and maximize widgets, as well as rearrange the widgets.

For more information, see:

- Understanding Widget Preferences, page 40-1
- Understanding the Predefined Widgets, page 40-2
- Working with the Dashboard, page 40-5

Understanding Widget Preferences

License: Any

Each widget has a set of preferences that determines its behavior.

Widget preferences can be simple. For example, you can set preferences for the Current Interface Status widget, which displays the current status of all enabled interfaces on the internal network. You can only configure the update frequency for this widget.

To modify a widget's preferences:

Step 1 On the title bar of the widget whose preferences you want to change, click the show preferences icon ( ).

The preferences section for that widget appears.

Step 2 Make changes as needed.

Your changes take effect immediately. For information on the preferences you can specify for individual widgets, see Understanding the Predefined Widgets, page 40-2.

Step 3 On the widget title bar, click the hide preferences icon ( ) to hide the preferences section.

Understanding the Predefined Widgets

License: Any

The ASA FirePOWER module is delivered with several predefined widgets that can provide you with at-a-glance views of current system status.

For detailed information on the widgets, see the following sections:

- Understanding the Appliance Information Widget, page 40-2
- Understanding the Current Interface Status Widget, page 40-3
- Understanding the Disk Usage Widget, page 40-3
- Understanding the Product Licensing Widget, page 40-4
- Understanding the Product Updates Widget, page 40-4
- Understanding the System Load Widget, page 40-5
- Understanding the System Time Widget, page 40-5

Understanding the Appliance Information Widget

License: Any

The Appliance Information widget provides:

- the name, IPv4 address, IPv6 address, and model of the appliance
- the versions of the ASA FirePOWER module software, rule update, vulnerability database (VDB), and geolocation update installed on the appliance.

You can configure the widget to display more or less information by modifying the widget preferences to display a simple or an advanced view; the preferences also control how often the widget updates. For more information, see Understanding Widget Preferences, page 40-1.

Understanding the Current Interface Status Widget

License: Any

The Current Interface Status widget shows the status of all interfaces on the appliance, enabled or unused. For each interface, the widget provides:

- the name of the interface
- the link state of the interface
- the link mode (for example, 100Mb full duplex, or 10Mb half duplex) of the interface
- the type of interface, that is, copper or fiber
- the amount of data received (Rx) and transmitted (Tx) by the interface

The color of the ball representing link state indicates the current status, as follows:

- green: link is up and at full speed
- yellow: link is up but not at full speed
- red: link is not up
- gray: link is administratively disabled
- blue: link state information is not available (for example, ASA)

The widget preferences control how often the widget updates. For more information, see Understanding Widget Preferences, page 40-1.

Understanding the Disk Usage Widget

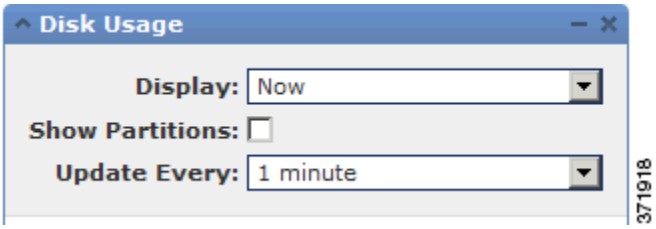
License: Any

The Disk Usage widget displays the space used on the hard drive, based on disk usage category. It also indicates the space used on and capacity of each partition of the appliance's hard drive. The By Category stacked bar displays each disk usage category as a proportion of the total available disk space used. The following table describes the available categories.

Table 40-1 Disk Usage Categories

Disk Usage Category	Description
Events	all events logged by the system
Files	all files stored by the system
Backups	all backup files
Updates	all files related to updates, such as rule updates and system updates
Other	system troubleshooting files and other miscellaneous files
Free	free space remaining on the appliance

You can configure the widget to display only the By Category stacked bar, or you can show the stacked bar plus the admin (`/`), `/Volume`, and `/boot` partition usage, as well as the `/var/storage` partition if the malware storage pack is installed, by modifying the widget preferences.



The widget preferences also control how often the widget updates, as well as whether it displays the current disk usage or collected disk usage statistics over the dashboard time range. For more information, see Understanding Widget Preferences, page 40-1.

Understanding the Product Licensing Widget

License: Any

The Product Licensing widget shows the device and feature licenses currently installed. It also indicates the number of items (such as hosts or users) licensed and the number of remaining licensed items allowed.

The top section of the widget displays all device and feature licenses installed, including temporary licenses, while the Expiring Licenses section displays only temporary and expired licenses.

The bars in the widget background show the percentage of each type of license that is being used; you should read the bars from right to left. Expired licenses are marked with a strikethrough.

You can configure the widget to display either the features that are currently licensed, or all the features that you can license, by modifying the widget preferences. The preferences also control how often the widget updates. For more information, see Understanding Widget Preferences, page 40-1.

You can click any of the license types to go to the License page of the local configuration and add or delete feature licenses. For more information, see Licensing the ASA FirePOWER Module, page 45-1.

Understanding the Product Updates Widget

License: Any

The Product Updates widget provides you with a summary of the software (ASA FirePOWER module software and rule updates) currently installed on the appliance as well as information on available updates that you have downloaded, but not yet installed, for that software.

Note that the widget displays `Unknown` as the latest version of the software unless you have configured a scheduled task to download, push, or install software updates; the widget uses scheduled tasks to determine the latest version. For more information, see Scheduling Tasks, page 42-1.

The widget also provides you with links to pages where you can update the software.

You can configure the widget to hide the latest versions by modifying the widget preferences. The preferences also control how often the widget updates. For more information, see Understanding Widget Preferences, page 40-1.

On the Product Updates widget, you can:

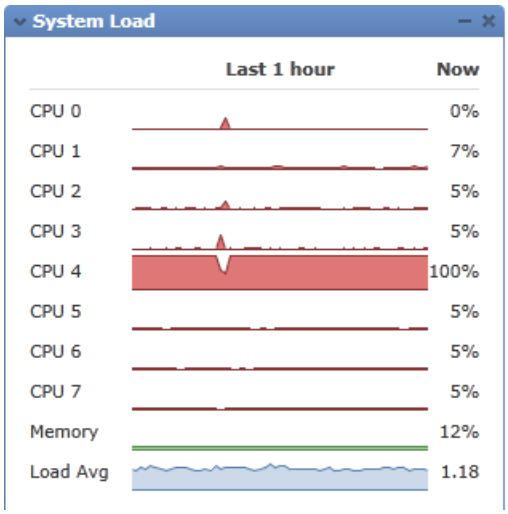
- manually update an appliance by clicking the current version of the ASA FirePOWER module software, rule update, VDB, or geolocation update:
- to update the system software, VDB, or geolocation database, see Updating ASA FirePOWER Module Software, page 46-1.

- to import the newest rule update, see Importing Rule Updates and Local Rule Files, page 46-9.
- create a scheduled task to download the latest version of the ASA FirePOWER module software, VDB, or rule update by clicking the latest version; see Scheduling Tasks, page 42-1.

Understanding the System Load Widget

License: Any

The System Load widget shows the CPU usage (for each CPU), memory (RAM) usage, and system load (also called the load average, measured by the number of processes waiting to execute) on the appliance, both currently and over the dashboard time range.

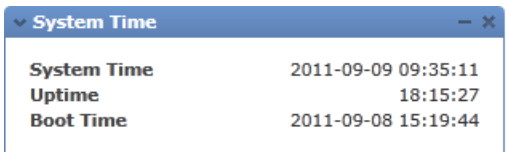


You can configure the widget to show or hide the load average by modifying the widget preferences. The preferences also control how often the widget updates. For more information, see Understanding Widget Preferences, page 40-1.

Understanding the System Time Widget

License: Any

The System Time widget shows the local system time, uptime, and boot time for the appliance.



You can configure the widget to hide the boot time by modifying the widget preferences. The preferences also control how often the widget synchronizes with the appliance's clock. For more information, see Understanding Widget Preferences, page 40-1.

Working with the Dashboard

License: Any

You can view and modify the widgets that appear on the dashboard.

For more information on working with the dashboard, see:

- Viewing the Dashboard, page 40-6
- Modifying the Dashboard, page 40-6
- Exporting Configurations, page B-1

Viewing the Dashboard

License: Any

At any time, to view the dashboard for your ASA FirePOWER module, select **Home > ASA FirePOWER Dashboard**.

The dashboard has a time range that constrains its widgets. You can change the time range to reflect a period as short as the last hour (the default) or as long as the last year. When you change the time range, the widgets that can be constrained by time automatically update to reflect the new time range.

Note that not all widgets can be constrained by time. For example, the dashboard time range has no effect on the Appliance Information widget, which provides information that includes the appliance name, model, and current version of the ASA FirePOWER module software.

To view the dashboard:

Step 1 Select **Home > ASA FirePOWER Dashboard**.

The ASA FirePOWER dashboard appears.

To change the dashboard time range:

Step 1 From the **Show the Last** drop-down list, choose a dashboard time range.

All appropriate widgets on the page update to reflect the new time range.

Modifying the Dashboard

License: Any

The dashboard displays widgets in a three-column layout. You can minimize and maximize widgets, as well as rearrange the widgets.

For more information, see the following sections:

- Rearranging Widgets, page 40-6
- Minimizing and Maximizing Widgets, page 40-7

Rearranging Widgets

License: Any

You can change the location of any widget.

To move a widget:

Step 1 Click the title bar of the widget you want to move, then drag it to its new location.

Minimizing and Maximizing Widgets

License: Any

You can minimize widgets to simplify your view, then maximize them when you want to see them again.

To minimize a widget:

Step 1 Click the minimize icon (−) in a widget's title bar.

To maximize a widget:

Step 1 Click the maximize icon (□) in a minimized widget's title bar.



Using ASA FirePOWER Reporting

You can view reports on various time periods to analyze the traffic on your network. Reports aggregate information on various aspects of your network traffic. In most cases, you can drill down from general information to specific information. For example, you can view a report on all users, then view details about specific users.

Overview and detail reports include multiple report components such as top policies and web categories. These reports show the most often occurring items of that type for the report you are viewing. For example, if you are viewing the detail report for a specific user, the top policies show the policy hits most associated with that user.

For more information, see:

- Understanding Available Reports, page 41-1
- Report Basics, page 41-2

Understanding Available Reports

License: Any

Available reports include the main reports available in the ASA FirePOWER module. You can view these reports from the ASA FirePOWER Reporting menu.

In general, you can click on many items, including names and View More links, to get more detailed information about individual items or about the monitored category as a whole.

Network Overview

This report shows summary information about the traffic in the network. Use this information to help identify areas that need deeper analysis, or to verify that the network is behaving within general expectations.

Users

This report shows the top users of your network. Use this information to help identify anomalous activity for a user.

Tip User names are available only when user identity information is associated with traffic flows. If you want to ensure that user identity is available in reports for the majority of traffic, the access control policy should use active authentication.

Applications

This report displays applications, which represent the content or requested URL for HTTP traffic detected in the traffic that triggered an intrusion event. Note that if the module detects an application protocol of HTTP, but cannot detect a specific web application, the module supplies a generic web browsing designation here.

Web categories

This report shows which categories of web sites, such as gambling, advertisements, or search engines and portals are being used in the network based on the categorization of web sites visited. Use this information to help identify the top categories visited by users and to determine whether your access control policies are sufficiently blocking undesired categories.

Policies

This report shows how your access control policies have been applied to traffic in the network. Use this information to help evaluate policy efficacy.

Ingress zones

This report displays the ingress security zone of the packet that triggered an event.

Egress zones

This report displays the egress security zone of the packet that triggered the event.

Destinations

This report shows which applications, such as Facebook, are being used in the network based on the analysis of the traffic in the network. Use this information to help identify the top applications used in the network and to determine whether additional access control policies are needed to reduce the usage of unwanted applications.

Attackers

This report displays the source IP addresses, used by the sending hosts, that triggered an event.

Targets

This report displays the destination IP addresses, used by the receiving hosts, that triggered an event.

Threats

This report displays the unique identifying number and explanatory text assigned to each detected threat to your network.

Files logs

This report displays the type of files detected, for example, HTML or MSEXE.

Report Basics

License: Any

The following sections explain the basics of using reports. These topics apply to reports in general and not to any single specific report.

For more information, see:

- Understanding Report Data, page 41-3
- Drilling into Reports, page 41-3
- Changing the Report Time Range, page 41-3
- Controlling the Data Displayed in Reports, page 41-4
- Understanding Report Columns, page 41-5

Understanding Report Data

License: Any

Report data is collected immediately from the device, so there is little lag time between the data reflected in a report and network activity. However, keep the following points in mind when analyzing the data:

- Data is collected for traffic that matches an access control policy applied to your ASA FirePOWER module.
- Data is aggregated into 5 minute buckets, and 30 minute and one hour graphs show data points in 5 minute increments. At the end of the hour, the 5 minute buckets are aggregated into one hour buckets, which are subsequently aggregated into day and week buckets. The 5 minute buckets are kept for 7 days, the one hour buckets for 31 days, and the day buckets for up to 365 days. The farther back you look, the more aggregated the data. When you query for old data, you get the best results if you align your queries to the availability of these data buckets.



Note If a data point is missing, for example, because the device was unreachable for longer than 5 minutes, there will be gaps in line charts.

Drilling into Reports

License: Any

Reports include many links to help you drill down to the information that you need. Mouse over items to see which ones might take you to more information about the item.

For example, in a typical reporting item, you can click the View More link to go to the summary report for that item.

You can also get to a detail report on a specific item by clicking the item in a summary report. For example, clicking Hypertext Transfer Protocol (HTTP) in the applications summary report takes you to the applications detail report for HTTP.

Changing the Report Time Range

License: Any

When you view a report, you can change the time range that defines the information to include in the report using the Time Range list. The time range list appears at the top of each report, and allows you to select predefined time ranges, such as the last hour or week, or to define a custom time range with specific start and end times. The time range you select is carried over to any other report that you view until you change the selection.

Reports automatically update every 10 minutes.

The following table explains the time range options.

Table 41-1 Time Ranges for reports

Time Range	Data Returned In
Last 30 minutes	30 complete minutes in five minute intervals, plus up to five additional minutes.
Last hour	60 complete minutes in five minute intervals, plus up to five additional minutes.
Last 24 hours	One hour intervals for the last 24 hours rounded to the previous hour boundary. For example, if the current time is 13:45, the Last 24 Hour period is from 13:00 yesterday to 13:00 today.
Last 7 days	One hour intervals for the last seven days rounded to the previous hour boundary.
Last 30 days	One day intervals for the last 30 days starting from the previous midnight.
Custom Range	The time range you define. Edit boxes are displayed for start date, start time, end date, and end time; click in each box and select the desired value. Click **Apply** to update the report when you are finished. When constructing a custom time range, you should align your range with the availability of data buckets. For ranges 7-31 days in the past, align your query on the hour. For older ranges, align them on the day; for ranges over a year, align them on the week.

Controlling the Data Displayed in Reports

License: Any

Overview and detail reports include several subordinate reports such as Top Policies and Web Categories. Each report panel includes controls that let you view different aspects of the data. You can use the following controls:

Transactions or Data Usage

Click these links to view charts based on the number of transactions or the amount of data in the transactions.

All, Denied, Allowed

The unlabeled drop-down list in the upper right of each report includes these options. Use them to change whether you see denied connections only, allowed connections only, or all connections whether denied or allowed.

View More

Click the View More link to go to the report for the item you are viewing. For example, clicking **View More** in the Web Categories chart of the Destinations report takes you to the Web Categories report. If you are viewing the report in a detailed report, you go to the detailed Web Categories report for the item you are viewing details about.

Understanding Report Columns

License: Any

Reports typically contain one or more tables to present information in addition to the information displayed in graphical format.

- The meaning of many columns is modified by the report in which they are included. For example, the transactions column shows the number of transactions for the type of item reported on. You can also toggle the values between raw numbers and as a percentage of the total reported raw values for the item by clicking **Values** or **Percentages**.
- You can change the sort order of the columns by clicking the column heading.

The following table explains the standard columns that you can find in the various reports.

Table 41-2 Report Columns

Column	Description
Transactions	The total number of transactions for the reported item.
Transactions allowed	The number of transactions that were allowed for the reported item.
Transactions denied	The number of transactions that were blocked (based on policy) for the reported item.
Total bytes	The sum of bytes sent and received for the reported item.
Bytes received	The number of bytes received for the reported item.
Total Bytes Sent	The number of bytes sent for the reported item.



Scheduling Tasks

You can schedule many different types of administrative tasks to run at designated times, either once or on a recurring basis.

Note Some tasks (such as those involving automated software updates) may place a significant load on networks with low bandwidths. You should schedule tasks like these to run during periods of low network use.

See the following sections for more information:

- Configuring a Recurring Task, page 42-1 explains how to set up a scheduled task so that it runs at regular intervals.
- Automating Backup Jobs, page 42-2 provides procedures for scheduling backup jobs.
- Automating Applying an Intrusion Policy, page 42-3 provides procedures for queuing an intrusion policy apply.
- Automating Geolocation Database Updates, page 42-4 provides procedures for scheduling automatic updates of the geolocation database (GeoDB).
- Automating Software Updates, page 42-5 provides procedures for scheduling the download, push, and installation of software updates.
- Automating URL Filtering Updates, page 42-7 provides procedures for automating updates of URL filtering data.
- Viewing Tasks, page 42-8 describes how to view and manage tasks after they are scheduled.
- Editing Scheduled Tasks, page 42-9 describes how to edit an existing task.
- Deleting Scheduled Tasks, page 42-10 describes how to delete one-time tasks and all instances of recurring tasks.

Configuring a Recurring Task

License: Any

You set the frequency for a recurring task using the same process for all types of tasks.

Note that the time displayed on most pages on the user interface is the local time, which is determined by using the time zone you specify in your local configuration. Further, the ASA FirePOWER module automatically adjusts its local time display for daylight saving time (DST), where appropriate. However, recurring tasks that span the transition dates from DST to standard time and back do not adjust for the

transition. That is, if you create a task scheduled for 2:00 AM during standard time, it will run at 3:00 AM during DST. Similarly, if you create a task scheduled for 2:00 AM during DST, it will run at 1:00 AM during standard time.

To configure a recurring task:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Tools > Scheduling**.

The Scheduling page appears.

Step 2 Click **Add Task**.

The New Task page appears.

Step 3 From the **Job Type** list, select the type of task that you want to schedule.

Each of the types of tasks you can schedule is explained in its own section.

Step 4 For the **Schedule task to run** option, select **Recurring**.

The page reloads with the recurring task options.

Step 5 In the **Start On** field, specify the date when you want to start your recurring task. You can use the drop-down list to select the month, day, and year.

Step 6 In the **Repeat Every** field, specify how often you want the task to recur. You can specify a number of hours, days, weeks, or months.



Tip You can either type a number or click the up icon (▲) and the down (▼) icon to specify the interval. For example, type `2` and select Days to run the task every two days.

Step 7 In the **Run At** field, specify the time when you want to start your recurring task.

Step 8 If you selected `Weeks` for **Repeat Every**, a **Repeat On** field appears. Select the check boxes next to the days of the week when you want to run the task.

Step 9 If you selected `Months` for **Repeat Every**, a **Repeat On** field appears. Use the drop-down list to select the day of the month when you want to run the task.

The remaining options on the New Task page are determined by the task you are creating. See the following sections for more information:

- Automating Backup Jobs, page 42-2
- Automating Applying an Intrusion Policy, page 42-3
- Automating Software Updates, page 42-5
- Automating URL Filtering Updates, page 42-7

Automating Backup Jobs

You can use the scheduler to automate backups of your ASA FirePOWER module. You must design a backup profile before you can configure a backup as a scheduled task. For more information, see Creating Backup Profiles, page 48-3.

To automate backup tasks:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Tools > Scheduling**.

The Scheduling page appears.

Step 2 Click **Add Task**.

The New Task page appears.

Step 3 From the **Job Type** list, select **Backup**.

The page reloads to show the backup options.

Step 4 Specify how you want to schedule the backup, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time. The **Current Time** field indicates the current time on the appliance.
- For recurring tasks, you have several options for setting the interval between instances of the task. See Configuring a Recurring Task, page 42-1 for details.

Step 5 In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.

Step 6 From the **Backup Profile** list, select the appropriate backup profile.

For more information on creating new backup profiles, see Creating Backup Profiles, page 48-3.

Step 7 Optionally, in the **Comment** field, type a comment using up to 255 alphanumeric characters, spaces, or periods.

Tip The comment field appears in the View Tasks section of the page, so you should try to keep it relatively short.

Step 8 Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want task status messages sent.

You must have a valid email relay server configured to send status messages. See Configuring a Mail Relay Host and Notification Address, page 43-6 for more information about configuring a relay host.

Step 9 Click **Save**.

The task is added. You can check the status of a running task on the Task Status page; see Viewing the Status of Long-Running Tasks, page C-1.

Automating Applying an Intrusion Policy

License: Protection

You can queue an intrusion policy apply to the ASA FirePOWER module. This task only applies the intrusion policy if an access control policy that references the intrusion policy is applied to the ASA FirePOWER module when the task runs. Otherwise, the task aborts before completion.

You must associate an intrusion policy with an access control policy and apply the access control policy to a device before scheduling this task; see Controlling Traffic Using Intrusion and File Policies, page 11-1.

To queue a policy apply:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Tools > Scheduling**.

The schedule calendar page for the current month appears.

Step 2 Click **Add Task**.

The New Task page appears.

Step 3 From the **Job Type** list, select **Queue Intrusion Policy Apply**.

The page reloads to show the options for queuing a policy apply.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time. The **Current Time** field indicates the current time on the ASA FirePOWER module.
- For recurring tasks, you have several options for setting the interval between instances of the task. See Configuring a Recurring Task, page 42-1 for details.

Step 5 In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.

Step 6 In the **Intrusion Policy** field, you have the following options:

- Select an intrusion policy to apply to the ASA FirePOWER module.
- Select **All intrusion policies** to apply all intrusion policies already applied to the ASA FirePOWER module.

Step 7 Optionally, in the **Comment** field, type a comment using up to 255 alphanumeric characters, spaces, or periods.

Tip The comment field appears in the Tasks Details section at the bottom of the schedule calendar page, so you should limit the size of your comment.

Step 8 Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want task status messages sent.

You must have a valid email relay server configured to send status messages. See Configuring a Mail Relay Host and Notification Address, page 43-6 for more information about configuring a relay host.

Step 9 Click **Save**.

The task is added. You can check the status of a running task in the Task Details section of the calendar page; see Viewing the Status of Long-Running Tasks, page C-1.

Step 10 To edit your saved task, click the task anywhere it appears on the schedule calendar page.

The Task Details section appears at the bottom of the page. To make any changes, click the edit icon ( ).

Automating Geolocation Database Updates

License: Any

You can use the scheduler to automate recurring geolocation database (GeoDB) updates. Recurring GeoDB updates run once every 7 days (weekly); you can configure the time the update recurs each week. For more information on GeoDB updates, see Updating the Geolocation Database, page 46-19.

To automate geolocation database updates:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Updates**.

The Product Updates page appears.

Step 2 Click the **Geolocation Updates** tab.

The Geolocation Updates page appears.

Step 3 Under **Recurring Geolocation Updates**, select the **Enable Recurring Weekly Updates** check box.

The Update Start Time field appears.

Step 4 In the **Update Start Time** field, specify the time and day of the week when you want weekly GeoDB updates to occur.

Step 5 Click **Save**.

The task is added. You can check the status of a running task on the Task Status page; see Viewing the Status of Long-Running Tasks, page C-1.

Automating Software Updates

You can automatically download and apply most patches and feature releases to the ASA FirePOWER module.

Note You must manually upload and install updates in two situations. First, you cannot schedule major updates to the ASA FirePOWER module. Second, you cannot schedule updates for or pushes from appliances that cannot access the Support Site. For information on manually updating the ASA FirePOWER module, see Updating ASA FirePOWER Module Software, page 46-1.

If you want to have more control over this process, you can use the **Once** option to download and install updates during off-peak hours after you learn that an update has been released.

See the following sections for more information:

- Automating Software Downloads, page 42-5
- Automating Software Installs, page 42-6

Automating Software Downloads

You can create a scheduled task that automatically downloads the latest software updates from Cisco. You can use this task to schedule download of updates you plan to install manually.

To automate software update downloads:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Tools > Scheduling**.

The Scheduling page appears.

Step 2 Click **Add Task**.

The New Task page appears.

Step 3 From the **Job Type** list, select **Download Latest Update**.

The New Task page reloads to show the update options.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time. The **Current Time** field indicates the current time on the appliance.
- For recurring tasks, you have several options for setting the interval between instances of the task. See Configuring a Recurring Task, page 42-1 for details.

Step 5 In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.

Step 6 In the **Update Items** section, select **Software**.

Step 7 Optionally, in the **Comment** field, type a comment using up to 255 alphanumeric characters, spaces, or periods.

Tip The comment field appears in the View Tasks section of the page, so you should try to keep it relatively short.

Step 8 Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want task status messages sent.

You must have a valid email relay server configured to send status messages. See Configuring a Mail Relay Host and Notification Address, page 43-6 for more information about configuring a relay host.

Step 9 Click **Save**.

The task is added. You can check the status of a running task on the Task Status page; see Viewing the Status of Long-Running Tasks, page C-1.

Automating Software Installs

Caution Depending on the update being installed, the appliance may reboot after the software is installed.

To schedule a software installation task:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Tools > Scheduling**.

The Scheduling page appears.

Step 2 Click **Add Task**.

The New Task page appears.

Step 3 From the **Job Type** list, select **Install Latest Update**.

The page reloads to show the options for installing updates.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time. The **Current Time** field indicates the current time on the appliance.
- For recurring tasks, you have several options for setting the interval between instances of the task. See Configuring a Recurring Task, page 42-1 for details.

Step 5 In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.

Step 6 Optionally, in the **Comment** field, type a comment using up to 255 alphanumeric characters, spaces, or periods.

Tip The comment field appears in the View Tasks section of the page, so you should try to keep it relatively short.

Step 7 Optionally, in the **Email Status To:** field, type the email address (or multiple email addresses separated by commas) where you want task status messages sent.

You must have a valid email relay server configured to send status messages. See Configuring a Mail Relay Host and Notification Address, page 43-6 for more information about configuring a relay host.

Step 8 Click **Save**.

The task is added. You can check the status of a running task on the Task Status page; see Viewing the Status of Long-Running Tasks, page C-1.

Automating URL Filtering Updates

License: URL Filtering

You can use the scheduler to automate updates of URL filtering data from the Collective Security Intelligence Cloud. For a URL filtering update task to succeed:

- The ASA FirePOWER module must have access to the Internet or it cannot contact the cloud.
- You must enable URL filtering, as described in Enabling Cloud Communications, page 44-2.

Note that when you enable URL filtering, you can also enable automatic updates. This forces the ASA FirePOWER module to contact the cloud every 30 minutes for URL filtering data updates. If you have enabled automatic updates, you should **not** create a scheduled task to update URL filtering data.

Although daily updates tend to be small, if it has been more than five days since your last update, new URL filtering data may take up to 20 minutes to download, depending on your bandwidth. Then, it may take up to 30 minutes to perform the update itself.

To automate URL filtering data tasks:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Tools > Scheduling**.

The Scheduling page appears.

Step 2 Click **Add Task**.

The New Task page appears.

Step 3 From the **Job Type** list, select **Update URL Filtering Database**.

The page reloads to show the URL filtering update options.

Step 4 Specify how you want to schedule the update, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time. The **Current Time** field indicates the current time on the appliance.
- For recurring tasks, you have several options for setting the interval between instances of the task. See Configuring a Recurring Task, page 42-1 for details.

Step 5 In the **Job Name** field, type a name using up to 255 alphanumeric characters, spaces, or dashes.

Step 6 Optionally, in the **Comment** field, type a comment using up to 255 alphanumeric characters, spaces, or periods.

Tip The comment field appears in the View Tasks section of the page, so you should try to keep it relatively short.

Step 7 Optionally, in the **Email Status To** field, type the email address (or multiple email addresses separated by commas) where you want task status messages sent.

You must have a valid email relay server configured to send status messages. See Configuring a Mail Relay Host and Notification Address, page 43-6 for more information about configuring a relay host.

Step 8 Click **Save**.

The task is added. You can check the status of a running task on the Task Status page; see Viewing the Status of Long-Running Tasks, page C-1.

Viewing Tasks

After adding scheduled tasks, you can view them and evaluate their status. The View Options section of the page allows you to view scheduled tasks using a calendar and a list of scheduled tasks.

See the following sections for more information:

- Using the Calendar, page 42-8
- Using the Task List, page 42-9

Using the Calendar

The Calendar view option allows you to view which scheduled tasks occur on which day.

To view scheduled tasks using the calendar:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Tools > Scheduling**.

The Scheduling page appears.

Step 2 You can perform the following tasks using the calendar view:

- Click the double left arrow icon («) to move back one year.
- Click the single left arrow icon (‹) to move back one month.
- Click the single right arrow icon (›) to move forward one month.

- Click the double right arrow icon ( ) to move forward one year.
- Click **Today** to return to the current month and year.
- Click **Add Task** to schedule a new task.
- Click a date to view all scheduled tasks for the specific date in a task list table below the calendar.
- Click a specific task on a date to view the task in a task list table below the calendar.

Note For more information about using the task list, see Using the Task List.

Using the Task List

The Task List shows a list of tasks along with their status. The task list appears below the calendar when you open the calendar. In addition, you can access it by selecting a date or task from the calendar. See Using the Calendar, page 42-8 for more information.

Table 42-1 Task List Columns

Column	Description
Name	Displays the name of the scheduled task and the comment associated with it.
Type	Displays the type of scheduled task.
Start Time	Displays the scheduled start date and time.
Frequency	Displays how often the task is run.
Status	Describes the current status for a scheduled task: • A check mark icon ( ) indicates that the task ran successfully. • A question mark icon ( ) indicates that the task is in an unknown state. • An exclamation mark icon ( ) indicates that the task failed.
Creator	Displays the name of the user that created the scheduled task.
Edit	Edits the scheduled task.
Delete	Deletes the scheduled task.

Editing Scheduled Tasks

You can edit a scheduled task that you previously created. This feature is especially useful if you want to test a scheduled task once to make sure that the parameters are correct. Later, after the task completes successfully, you can change it to a recurring task.

To edit an existing scheduled task:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Tools > Scheduling**.

The Scheduling page appears.

Step 2 Click either the task that you want to edit or the day on which the task appears.

The Task Details table containing the selected task or tasks appears.

Step 3 Locate the task you want to edit in the table and click the edit icon ( )

The Edit Task page appears, showing the details of the task you selected.

Step 4 Edit the task to meet your needs, including the start time, the job name, the comment, and how often the task runs, once or recurring. You cannot change the type of job.

The remaining options are determined by the task you are editing. See the following sections for more information:

- Automating Backup Jobs, page 42-2
- Automating Software Updates, page 42-5
- Automating URL Filtering Updates, page 42-7

Step 5 Click **Save** to save your edits.

Your change are saved and the Scheduling page appears again.

Deleting Scheduled Tasks

There are two types of deletions you can perform from the Schedule View page. You can delete a specific one-time task that has not yet run or you can delete every instance of a recurring task. If you delete an instance of a recurring task, all instances of the task are deleted. If you delete a task that is scheduled to run once, only that task is deleted.

The following sections describe how to delete tasks:

- To delete all instances of a task, see Deleting a Recurring Task, page 42-10.
- To delete a single instance of a task, see Deleting a One-Time Task, page 42-11.

Deleting a Recurring Task

When you delete one instance of a recurring task, you automatically delete all instances of that task.

To delete a recurring task:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Tools > Scheduling**.

The Scheduling page appears.

Step 2 On the calendar, select an instance of the recurring task you want to delete.

The page reloads to display a table of tasks below the calendar.

Step 3 Locate an instance of the recurring task you want to delete in the table and click the delete icon ( ).

All instances of the recurring task are deleted.

Deleting a One-Time Task

You can delete a one-time scheduled task or delete the record of a previously run scheduled task using the task list.

To delete a single task or, if it has already run, delete a task record:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Tools > Scheduling**.

The Scheduling page appears.

Step 2 Click the task that you want to delete or the day on which the task appears.

A table containing the selected task or tasks appears.

Step 3 Locate the task you want to delete in the table and click the delete icon ( ).

The instance of the task you selected is deleted.



Managing System Policies

A system policy allows you to manage the following on your ASA FirePOWER module:

- audit log settings
- the mail relay host and notification address
- SNMP polling settings
- STIG compliance

See the following sections for more information:

- Creating a System Policy, page 43-1
- Editing a System Policy, page 43-2
- Applying a System Policy, page 43-2
- Deleting System Policies, page 43-3

Creating a System Policy

License: Any

When you create a system policy, you assign it a name and a description. Next, you configure the various aspects of the policy, each of which is described in its own section.

Instead of creating a new policy, you can export a system policy from another ASA FirePOWER module and then import it onto your ASA FirePOWER module. You can then edit the imported policy to suit your needs before you apply it. For more information, see Importing and Exporting Configurations, page B-1.

To create a system policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > System Policy**.

The System Policy page appears.

Step 2 Click **Create Policy**.

The Create Policy page appears.

Step 3 From the drop-down list, select an existing policy to use as a template for your new system policy.

Step 4 Type a name for your new policy in the **New Policy Name** field.

Step 5 Type a description for your new policy in the **New Policy Description** field.

Step 6 Click **Create**.

Your system policy is saved and the Edit System Policy page appears. For information about configuring each aspect of the system policy, see one of the following sections:

- Configuring Audit Log Settings, page 43-5
- Configuring a Mail Relay Host and Notification Address, page 43-6
- Configuring SNMP Polling, page 43-8
- Enabling STIG Compliance, page 43-9

Editing a System Policy

License: Any

You can edit an existing system policy. If you edit a system policy that is currently applied to an ASA FirePOWER module, reapply the policy after you have saved your changes. For more information, see Applying a System Policy, page 43-2.

To edit an existing system policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > System Policy**.

The System Policy page appears, including a list of the existing system policies.

Step 2 Click the edit icon ( ) next to the system policy that you want to edit.

The Edit Policy page appears. You can change the policy name and policy description. For information about configuring each aspect of the system policy, see one of the following sections:

- Configuring Audit Log Settings, page 43-5
- Configuring a Mail Relay Host and Notification Address, page 43-6
- Configuring SNMP Polling, page 43-8
- Enabling STIG Compliance, page 43-9

Note If you are editing a system policy applied to an ASA FirePOWER module, make sure you reapply the updated policy when you are finished. See Applying a System Policy, page 43-2.

Step 3 Click **Save Policy and Exit** to save your changes. The changes are saved, and the System Policy page appears.

Applying a System Policy

License: Any

You can apply a system policy to an ASA FirePOWER module. If a system policy is already applied, any changes you make do not take effect until you reapply it.

To apply a system policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > System Policy**.

The System Policy page appears.

Step 2 Click the apply icon ( ) next to the system policy that you want to apply.

Step 3 Click **Apply**.

The System Policy page appears. A message indicates the status of applying the system policy.

Deleting System Policies

License: Any

You can delete a system policy, even if it is in use. If the policy is still in use, it is used until a new policy is applied. Default system policies cannot be deleted.

To delete a system policy:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > System Policy**.

The System Policy page appears.

Step 2 Click the delete icon ( ) next to the system policy that you want to delete. To delete the policy, click **OK.**

The System Policy page appears. A pop-up message appears, confirming the policy deletion.

Configuring a System Policy

License: Any

You can configure various system policy settings. For information about configuring each aspect of the system policy, see one of the following sections:

- Configuring the Access List for Your Appliance, page 43-3
- Configuring Audit Log Settings, page 43-5
- Configuring a Mail Relay Host and Notification Address, page 43-6
- Configuring SNMP Polling, page 43-8
- Enabling STIG Compliance, page 43-9

Configuring the Access List for Your Appliance

License: Any

The Access List page allows you to control which computers can access your appliance on specific ports. By default, port 443 (Hypertext Transfer Protocol Secure, or HTTPS), which is used to access the web interface, and port 22 (Secure Shell, or SSH), which is used to access the command line, are enabled for any IP address. You can also add SNMP access over port 161. Note that you must add SNMP access for any computer you plan to use to poll for SNMP information.

Caution By default, access to the appliance is **not** restricted. To operate the appliance in a more secure environment, consider adding access to the appliance for specific IP addresses and then deleting the default `any` option.

The access list is part of the system policy. You can specify the access list either by creating a new system policy or by editing an existing system policy. In either case, the access list does not take effect until you apply the system policy.

Note that this access list does not also control external database access. For more information on the external database access list, see Enabling Cloud Communications, page 44-2.

To configure the access list:

Access: Admin

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > System Policy**.

The System Policy page appears.

Step 2 You have the following options:

- To modify the access list in an existing system policy, click the edit icon ( ) next to the system policy.
- To configure the access list as part of a new system policy, click **Create Policy**.

  Provide a name and description for the system policy as described in Creating a System Policy, page 43-1, and click **Save**.

In either case, the Access List page appears.

Step 3 Optionally, to delete one of the current settings, click the delete icon ( ).

The setting is removed.

Caution If you delete access for the IP address that you are currently using to connect to the appliance interface, and there is no entry for "`IP=any port=443`", you will lose access to the system when you apply the policy.

Step 4 Optionally, to add access for one or more IP addresses, click **Add Rules**.

The Add IP Address page appears.

Step 5 In the **IP Address** field, you have the following options, depending on the IP addresses you want to add:

- an exact IP address (for example, 192.168.1.101)
- an IP address block using CIDR notation (for example, 192.168.1.1/24)

  For information on using CIDR in the Firepower system, see IP Address Conventions, page 1-4.
- `any`, to designate any IP address

Step 6 Select **SSH**, **HTTPS**, **SNMP**, or a combination of these options to specify which ports you want to enable for these IP addresses.

Step 7 Click **Add**.

The Access List page appears again, reflecting the changes you made.

Step 8 Click **Save Policy and Exit**.

The system policy is updated. Your changes do not take effect until you apply the system policy. See Applying a System Policy, page 43-2 for more information.

Configuring Audit Log Settings

License: Any

You can configure the system policy so that the ASA FirePOWER module streams an audit log to an external host.

Note You must ensure that the external host is functional and accessible from the ASA FirePOWER module sending the audit log.

The sending host name is part of the information sent. You can further identify the audit log stream with a facility, a severity, and an optional tag. The ASA FirePOWER module does not send the audit log until you apply the system policy.

After you apply a policy with this feature enabled, and your destination host is configured to accept the audit log, the syslog messages are sent. The following is an example of the output structure:

```
Date Time Host [Tag] Sender: [User_Name]@[User_IP], [Subsystem], [Action]
```

where the local date, time, and hostname precede the bracketed optional tag, and the sending device name precedes the audit log message.

For example:

```
Mar 01 14:45:24 localhost [TAG] Dev-DC3000: admin@10.1.1.2, Operations > Monitoring, Page
View
```

To configure the audit log settings:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > System Policy**.

The System Policy page appears.

Step 2 You have the following options:

- To modify the audit log settings in an existing system policy, click the edit icon ( ) next to the system policy.
- To configure the audit log settings as part of a new system policy, click **Create Policy**.

  Provide a name and description for the system policy as described in Creating a System Policy, page 43-1, and click **Save**.

Step 3 Click **Audit Log Settings**.

The Audit Log Settings page appears.

Step 4 Select **Enabled** from the **Send Audit Log to Syslog** drop-down menu. (The default setting is **Disabled**.)

Step 5 Designate the destination host for the audit information by using the IP address or the fully qualified name of the host in the **Host** field. The default port (514) is used.

Caution If the computer you configure to receive an audit log is not set up to accept remote messages, the host will not accept the audit log.

Step 6 Select a syslog facility from the **Facility** field.

Step 7 Select a severity from the **Severity** field.

Step 8 Optionally, insert a reference tag in the **Tag (optional)** field.

Step 9 To send regular audit log updates to an external HTTP server, select **Enabled** from the **Send Audit Log to HTTP Server** drop-down list. The default setting is **Disabled**.

Step 10 In the **URL to Post Audit** field, designate the URL where you want to send audit information. You must enter an URL that corresponds to a listener program that expects the HTTP POST variables as listed:

- `subsystem`
- `actor`
- `event_type`
- `message`
- `action_source_ip`
- `action_destination_ip`
- `result`
- `time`
- `tag` (if defined, as above)

Caution To allow encrypted posts, you must use an HTTPS URL. Note that sending audit information to an external URL may affect system performance.

Step 11 Click **Save Policy and Exit**.

The system policy is updated. Your changes do not take effect until you apply the system policy. See Applying a System Policy, page 43-2 for more information.

Configuring a Mail Relay Host and Notification Address

License: Any

You must configure a mail host if you plan to:

- email event-based reports
- email status reports for scheduled tasks
- email change reconciliation reports
- email data pruning notifications
- use email for intrusion event alerting

You can select an encryption method for the communication between appliance and mail relay host, and can supply authentication credentials for the mail server if needed. After configuring settings, you can test the connection between the appliance and the mail server using the supplied settings.

To configure a mail relay host:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > System Policy**.

The System Policy page appears.

Step 2 You have the following options:

- To modify the email settings in an existing system policy, click the edit icon ( ) next to the system policy.
- To configure the email settings as part of a new system policy, click **Create Policy**.

  Provide a name and description for the system policy as described in Creating a System Policy, page 43-1, and click **Save**.

Step 3 Click **Email Notification**.

The Configure Email Notification page appears.

Step 4 In the **Mail Relay Host** field, type the hostname or IP address of the mail server you want to use.

Note The mail host you enter must allow access from the appliance.

Step 5 Enter the port number to use on the email server in the **Port Number** field. Typical ports include 25, when using no encryption, 465, when using SSLv3, and 587, when using TLS.

Step 6 To select an encryption method, you have the following options:

- To encrypt communications between the appliance and the mail server using Transport Layer Security, select **TLS** from the **Encryption Method** drop-down list.
- To encrypt communications between the appliance and the mail server using Secure Socket Layers, select **SSLv3** from the **Encryption Method** drop-down list.
- To allow unencrypted communication between the appliance and the mail server, select **None** from the **Encryption Method** drop-down list.

Note that certificate validation is not required for encrypted communication between the appliance and mail server.

Step 7 Enter a valid email address in the **From Address** field for use as the source email address for messages sent by the appliance.

Step 8 Optionally, to supply a user name and password when connecting to the mail server, select **Use Authentication**. Enter a user name in the **Username** field. Enter a password in the **Password** field.

Step 9 To send a test email using the configured mail server, click **Test Mail Server Settings**.

A message appears next to the button indicating the success or failure of the test.

Step 10 Click **Save Policy and Exit**.

The system policy is updated. Your changes do not take effect until you apply the system policy. See Applying a System Policy, page 43-2 for more information.

Configuring SNMP Polling

License: Any

You can enable Simple Network Management Protocol (SNMP) polling of an appliance using the system policy. The SNMP feature supports use of versions 1, 2, and 3 of the SNMP protocol.

Note that enabling the system policy SNMP feature does not cause the appliance to send SNMP traps; it only makes the information in the MIBs available for polling by your network management system.

Note You must add SNMP access for any computer you plan to use to poll the appliance. For more information, see Configuring the Access List for Your Appliance, page 43-3. Note that the SNMP MIB contains information that could be used to attack your appliance. Cisco recommends that you restrict your access list for SNMP access to the specific hosts that will be used to poll for the MIB. Cisco also recommends you use SNMPv3 and use strong passwords for network management access.

To configure SNMP polling:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > System Policy**.

The System Policy page appears.

Step 2 You have the following options:

- To modify the SNMP polling settings in an existing system policy, click the edit icon ( ) next to the system policy.
- To configure the SNMP polling settings as part of a new system policy, click **Create Policy**.

  Provide a name and description for the system policy as described in Creating a System Policy, page 43-1, and click **Create**.

Step 3 If you have not already added SNMP access for each computer you plan to use to poll the appliance, do so now. For more information, see Configuring the Access List for Your Appliance, page 43-3.

Step 4 Click **SNMP**.

The SNMP page appears.

Step 5 From the **SNMP Version** drop-down list, select the SNMP version you want to use.

The drop-down list displays the version you selected.

Step 6 You have the following options:

- If you selected **Version 1** or **Version 2**, type the SNMP community name in the **Community String** field. Go to step 15.
- If you selected **Version 3**, click **Add User** to display the user definition page.

Step 7 Enter a username in the **Username** field.

Step 8 Select the protocol you want to use for authentication from the **Authentication Protocol** drop-down list.

Step 9 Type the password required for authentication with the SNMP server in the **Authentication Password** field.

Step 10 Retype the authentication password in the **Verify Password** field just below the **Authentication Password** field.

Step 11 Select the privacy protocol you want to use from the **Privacy Protocol** list, or select **None** to not use a privacy protocol.

Step 12 Type the SNMP privacy key required by the SNMP server in the **Privacy Password** field.

Step 13 Retype the privacy password in the **Verify Password** field just below the **Privacy Password** field.

Step 14 Click **Add**.

The user is added. You can repeat steps 6 through 13 to add additional users. Click the delete icon ( ) to delete a user.

Step 15 Click **Save Policy and Exit**.

The system policy is updated. Your changes do not take effect until you apply the system policy. See Applying a System Policy, page 43-2 for more information.

Enabling STIG Compliance

License: Any

Organizations within the United States federal government sometimes need to comply with a series of security checklists set out in Security Technical Implementation Guides (STIGs). The STIG Compliance option enables settings intended to support compliance with specific requirements set out by the United States Department of Defense.

Enabling STIG compliance does not guarantee strict compliance to all applicable STIGs.

When you enable STIG compliance, password complexity and retention rules for local shell access accounts change. In addition, you cannot use `ssh` remote storage when in STIG compliance mode.

Note that applying a system policy with STIG compliance enabled forces appliances to reboot. If you apply a system policy with STIG enabled to an appliance that already has STIG enabled, the appliance does not reboot. If you apply a system policy with STIG disabled to an appliance that has STIG enabled, STIG remains enabled and the appliance does not reboot.

Caution You cannot disable this setting without assistance from Support. In addition, this setting may substantially impact the performance of your system. Cisco does not recommend enabling STIG compliance except to comply with Department of Defense security requirements.

To enable STIG compliance:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > System Policy**.

The System Policy page appears.

Step 2 You have the following options:

- To modify the time settings in an existing system policy, click the edit icon ( ) next to the system policy.
- To configure the time settings as part of a new system policy, click **Create Policy**.

  Provide a name and description for the system policy as described in Creating a System Policy, page 43-1, and click **Save**.

Step 3 Click **STIG Compliance**.

The STIG Compliance page appears.

Step 4 If you want to *permanently* enable STIG compliance on the appliance, select **Enable STIG Compliance**.

Caution You cannot disable STIG compliance on an appliance after you apply a policy with STIG compliance enabled. If you need to disable compliance, contact Support.

Step 5 Click **Save Policy and Exit**.

The system policy is updated. Your changes do not take effect until you apply the system policy. See Applying a System Policy, page 43-2 for more information.

When you apply a system policy that enables STIG compliance to an appliance, note that the appliance reboots. Note that if you apply a system policy with STIG enabled to an appliance that already has STIG enabled, the appliance does not reboot.



Configuring ASA FirePOWER Module Settings

The following table summarizes an ASA FirePOWER module's local configuration.

Table 44-1 Local Configuration Options

Option	Description	For more information, see...
Information	Allows you to view current information about the appliance. You can also change the appliance name.	Viewing and Modifying the Appliance Information, page 44-1
Cisco CSI	Allows you to download URL filtering data from the Collective Security Intelligence Cloud, perform lookups for uncategorized URLs, and send diagnostic information on detected files to Cisco.	Enabling Cloud Communications, page 44-2

Viewing and Modifying the Appliance Information

License: Any

The Information page provides you with information about your ASA FirePOWER module. The information includes read-only information, such as the product name and model number, the operating system and version, and the current system policy. The page also provides you with an option to change the name of the appliance.

The following table describes each field.

Table 44-2 Appliance Information

Field	Description
Name	A name you assign to the appliance. Note that this name is only used within the context of the ASA FirePOWER module. Although you can use the hostname as the name of the appliance, entering a different name in this field does not change the hostname.
Product Model	The model name for the appliance.
Serial Number	The chassis serial number of the appliance.
Software Version	The version of the software currently installed.

Table 44-2 Appliance Information (continued)

Field	Description
Operating System	The operating system currently running on the appliance.
Operating System Version	The version of the operating system currently running on the appliance.
IPv4 Address	The IPv4 address of the default (`eth0`) management interface of the appliance. If IPv4 management is disabled for the appliance, this field indicates that.
IPv6 Address	The IPv6 address of the default (`eth0`) management interface of the appliance. If IPv6 management is disabled for the appliance, this field indicates that.
Current Policies	The appliance-level policies currently applied. If a policy has been updated since it was last applied, the name of the policy appears in italics.
Model Number	The model number for the appliance. This number may be important for troubleshooting.

To modify the appliance information:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > Configuration**.

The Information page appears.

Step 2 To change the appliance name, type a new name in the **Name** field.

The name **must** be alphanumeric characters and cannot be composed of numeric characters only.

Step 3 To save your changes, click **Save**.

The page refreshes and your changes are saved.

Enabling Cloud Communications

License: URL Filtering or Malware

The ASA FirePOWER module contacts Cisco's Collective Security Intelligence Cloud to obtain various types of information:

- File policies associated with access control rules allow devices to detect files transmitted in network traffic. The ASA FirePOWER module uses data from the Cisco cloud to determine if the files represent malware; see Understanding and Creating File Policies, page 35-4.
- When you enable URL filtering, the ASA FirePOWER module can retrieve category and reputation data for many commonly visited URLs, as well as perform lookups for uncategorized URLs. You can then quickly create URL conditions for access control rules; see Performing Reputation-Based URL Blocking, page 8-8.

Use the ASA FirePOWER module's local configuration to specify the following options:

Enable URL Filtering

You must enable this option to perform category and reputation-based URL filtering.

Query Cloud for Unknown URL

Allows the system to query the cloud when someone on your monitored network attempts to browse to a URL that is not in the local data set.

If the cloud does not know the category or reputation of a URL, or if the ASA FirePOWER module cannot contact the cloud, the URL does **not** match access control rules with category or reputation-based URL conditions. You cannot assign categories or reputations to URLs manually.

Disable this option if you do not want your uncategorized URLs to be cataloged by the Cisco cloud, for example, for privacy reasons.

Enable Automatic Updates

Allows the system to contact the cloud on a regular basis to obtain updates to the URL data in your appliances' local data sets. Although the cloud typically updates its data once per day, enabling automatic updates forces the ASA FirePOWER module to check every 30 minutes to make sure that you always have up-to-date information.

Although daily updates tend to be small, if it has been more than five days since your last update, new URL filtering data may take up to 20 minutes to download, depending on your bandwidth. Then, it may take up to 30 minutes to perform the update itself.

If you want to have strict control of when the system contacts the cloud, you can disable automatic updates and use the scheduler instead, as described in Automating URL Filtering Updates, page 42-7.

Note Cisco recommends that you either enable automatic updates or use the scheduler to schedule updates. Although you can manually perform on-demand updates, allowing the system to automatically contact the cloud on a regular basis provides you with the most up-to-date, relevant URL data.

Licensing

Performing category and reputation-based URL filtering and device-based malware detection require that you enable the appropriate licenses on your ASA FirePOWER module; see Licensing the ASA FirePOWER Module, page 45-1.

You **cannot** configure cloud connection options if you have no URL Filtering license on the ASA FirePOWER module. The Cisco CSI local configuration page displays only the options for which you are licensed. ASA FirePOWER modules with expired licenses cannot contact the cloud.

Note that, in addition to causing the URL Filtering configuration options to appear, adding a URL Filtering license to your ASA FirePOWER module automatically enables **Enable URL Filtering** and **Enable Automatic Updates**. You can manually disable the options if needed.

Internet Access

The system uses ports 80/HTTP and 443/HTTPS to contact the Cisco cloud.

The following procedures explain how to enable communications the Cisco cloud, and how to perform an on-demand update of URL data. Note that you cannot start an on-demand update if an update is already in progress.

To enable communications with the cloud:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Integration > Cisco CSI**.

The Information page appears.

Step 2 Click **Cisco CSI**.

The Cisco CSI page appears. If you have a URL Filtering license, the page displays the last time URL data was updated.

Step 3 Configure cloud connection options as described above.

You must **Enable URL Filtering** before you can **Enable Automatic Updates** or **Query Cloud for Unknown URLs**.

Step 4 Click **Save**.

Your settings are saved. If you enabled URL filtering, depending on how long it has been since URL filtering was last enabled, or if this is the first time you enabled URL filtering, the ASA FirePOWER module retrieves URL filtering data from the cloud.

To perform an on-demand update of the system's URL data:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Local > Configuration**.

The Information page appears.

Step 2 Click **URL Filtering**.

The URL Filtering page appears.

Step 3 Click **Update Now**.

The ASA FirePOWER module contacts the cloud and updates its URL filtering data if an update is available.

System Information

The System Information page displays information about your appliance, including the model, serial number, and version information for various components of the system. It also includes Cisco copyright information.

Time

You can view the current time and time source on the ASA FirePOWER module using the Time page.



Licensing the ASA FirePOWER Module

You can license a variety of features to create an optimal ASA FirePOWER deployment for your organization.

For more information, see:

- Understanding Licensing, page 45-1
- Viewing Your Licenses, page 45-4
- Adding a License to the ASA FirePOWER module, page 45-4
- Deleting a License, page 45-5

Understanding Licensing

License: Any

You can license a variety of features to create an optimal ASA FirePOWER deployment for your organization.

Licenses allow your device to perform a variety of functions including:

- intrusion detection and prevention
- Security Intelligence filtering
- file control and advanced malware protection
- application, user, and URL control

There are a few ways you may lose access to licensed features in the ASA FirePOWER module. You can remove licensed capabilities. Though there are some exceptions, you cannot use the features associated with an expired or deleted license.

This section describes the types of licenses available in an ASA FirePOWER module deployment. The licenses you can enable on an appliance can depend the other licenses enabled.

The following table summarizes ASA FirePOWER module licenses.

Table 45-1 ASA FirePOWER Module Licenses

License	Granted Capabilities	Requires
Protection	intrusion detection and prevention file control Security Intelligence filtering	none
Control	user and application control	Protection
Malware	advanced malware protection (network-based malware detection and blocking)	Protection
URL Filtering	category and reputation-based URL filtering	Protection

For more information, see:

- Protection, page 45-2
- Control, page 45-3
- Malware, page 45-3
- URL Filtering, page 45-3

Protection

License: Protection

A Protection license allows you to perform intrusion detection and prevention, file control, and Security Intelligence filtering:

- *Intrusion detection and prevention* allows you to analyze network traffic for intrusions and exploits and, optionally, drop offending packets.
- *File control* allows you to detect and, optionally, block users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols. With a Malware license (see Malware, page 45-3), you can also inspect and block a restricted set of those file types based on their malware dispositions.
- *Security Intelligence filtering* allows you to blacklist—deny traffic to and from—specific IP addresses, before the traffic is subjected to analysis by access control rules. Dynamic feeds allow you to immediately blacklist connections based on the latest intelligence. Optionally, you can use a "monitor-only" setting for Security Intelligence filtering.

Although you can configure an access control policy to perform Protection-related inspection without a license, you cannot apply the policy until you first add a Protection license to the ASA FirePOWER module.

If you delete your Protection license from the ASA FirePOWER module, the ASA FirePOWER module stops detecting intrusion and file events. Additionally, the ASA FirePOWER module will not contact the internet for either Cisco-provided or third-party Security Intelligence information. You cannot reapply existing policies until you re-enable Protection.

Because a Protection license is required for URL Filtering, Malware, and Control licenses, deleting or disabling a Protection license has the same effect as deleting or disabling your URL Filtering, Malware, or Control license.

Control

License: Control

A Control license allows you to implement user and application control by adding user and application conditions to access control rules. To enable Control, you must also enable Protection.

Although you can add user and application conditions to access control rules without a Control license, you cannot apply the policy until you first add a Control license to the ASA FirePOWER module.

If you delete your Control license, you cannot reapply existing access control policies if they include rules with user or application conditions.

URL Filtering

License: URL Filtering

URL filtering allows you to write access control rules that determine the traffic that can traverse your network based on URLs requested by monitored hosts, correlated with information about those URLs, which is obtained from the Cisco cloud by the ASA FirePOWER module. To enable URL Filtering, you must also enable a Protection license.

Tip Without a URL Filtering license, you can specify individual URLs or groups of URLs to allow or block. This gives you granular, custom control over web traffic, but does not allow you to use URL category and reputation data to filter network traffic.

URL filtering requires a subscription-based URL Filtering license. Although you can add category and reputation-based URL conditions to access control rules without a URL Filtering license, the ASA FirePOWER module will not contact the cloud for URL information. You cannot apply the access control policy until you first add a URL Filtering license to the ASA FirePOWER module.

You may lose access to URL filtering if you delete the license from the ASA FirePOWER module. Also, URL Filtering licenses may expire. If your license expires or if you delete it, access control rules with URL conditions immediately stop filtering URLs, and your ASA FirePOWER module can no longer contact the cloud. You cannot reapply existing access control policies if they include rules with category and reputation-based URL conditions.

Malware

License: Malware

A Malware license allows you to perform advanced malware protection, that is, use devices to detect and block malware in files transmitted over your network. To enable Malware on a device, you must also enable Protection.

You configure malware detection as part of a file policy, which you then associate with one or more access control rules. File policies can detect your users uploading or downloading files of specific types over specific application protocols. The Malware license allows you to inspect a restricted set of those file types for malware. The Malware license also allows you to add specific files to a file list and enable the file list within a file policy, allowing those files to be automatically allowed or blocked on detection.

Although you can add a malware-detecting file policy to an access control rule without a Malware license, the file policy is marked with a warning icon ( ) in the access control rule editor. Within the file policy, Malware Cloud Lookup rules are also marked with the warning icon. Before you can apply

an access control policy that includes a malware-detecting file policy, you **must** add a Malware license. If you later delete the license, you cannot reapply an existing access control policy to those devices if it includes file policies that perform malware detection.

If you delete your Malware license or it expires, the ASA FirePOWER module stops performing malware cloud lookups, and also stops acknowledging retrospective events sent from the Cisco cloud. You cannot reapply existing access control policies if they include file policies that perform malware detection. Note that for a very brief time after a Malware license expires or is deleted, the system can use cached dispositions for files detected by Malware Cloud Lookup file rules. After the time window expires, the system assigns a disposition of `Unavailable` to those files, rather than performing a lookup.

Viewing Your Licenses

License: Any

Use the Licenses page to view the licenses for an ASA FirePOWER module.

Other than the Licenses page, there are a few other ways you can view licenses and license limits:

- The Product Licensing dashboard widget provides an at-a-glance overview of your licenses.
- The Device page (**Configuration > ASA FirePOWER Configuration > Device Management > Device**) lists the licenses.

To view your licenses:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Licenses**.

The Licenses page appears.

Adding a License to the ASA FirePOWER module

License: Any

Before you add a license to the ASA FirePOWER module, make sure you have the activation key provided by Cisco when you purchased the license. You **must** add licenses before you can use licensed features.

Note If you add licenses after a backup has completed, these licenses will not be removed or overwritten if this backup is restored. To prevent a conflict on restore, remove those licenses before restoring the backup, noting where the licenses were used, and add and reconfigure them after restoring the backup. If a conflict occurs, contact Support.

To add a license:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Licenses**.

The Licenses page appears.

Step 2 Click **Add New License**.

The Add License page appears.

Step 3 Did you receive an email with your license?

- If yes, copy the license from the email, paste it into the **License** field, and click **Submit License**.

  If the license is correct, the license is added. Skip the rest of the procedure.

- If no, click **Get License**.

  The Product License Registration portal appears. If you cannot access the Internet, switch to a computer that can. Note the license key at the bottom of the page and browse to https://www.cisco.com/go/license.

Step 4 Follow the on-screen instructions to obtain your license, which will be sent to you in an email.

Tip You can also request a license on the **Licenses** tab after you log into the Support Site.

Step 5 Copy the license from the email, paste it into the **License** field in the ASA FirePOWER module's web user interface, and click **Submit License**.

If the license is valid, it is added.

Deleting a License

License: Any

Use the following procedure if you need to delete a license for any reason. Keep in mind that because Cisco generates licenses based on each ASA FirePOWER module's unique license key, you cannot delete a license from one ASA FirePOWER module and then reuse it on a different ASA FirePOWER module.

In most cases, deleting a license removes your ability to use features enabled by that license. For more information, see Understanding Licensing, page 45-1.

To delete a license:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Licenses**.

The Licenses page appears.

Step 2 Next to the license you want to delete, click the delete icon ( ).

Step 3 Confirm that you want to delete the license.

The license is deleted.



Updating ASA FirePOWER Module Software

Cisco electronically distributes several different types of updates, including major and minor updates to the ASA FirePOWER module software itself, as well as rule updates, geolocation database (GeoDB) updates, and Vulnerability Database (VDB) updates.

Caution This section contains general information on updating the ASA FirePOWER module. Before you update, including the VDB, GeoDB, or intrusion rules, you **must** read the release notes or advisory text that accompanies the update. The release notes provide important information, including prerequisites, warnings, and specific installation and uninstallation instructions.

Unless otherwise documented in the release notes or advisory text, updating does not modify configurations; the settings remain intact.

See the following sections for more information:

- Understanding Update Types, page 46-1
- Performing Software Updates, page 46-2
- Uninstalling Software Updates, page 46-7
- Updating the Vulnerability Database, page 46-8
- Importing Rule Updates and Local Rule Files, page 46-9
- Updating the Geolocation Database, page 46-19

Understanding Update Types

License: Any

Cisco electronically distributes several different types of updates, including major and minor updates to the ASA FirePOWER module software itself, as well as intrusion rule updates and VDB updates.

The following table describes the types of updates provided by Cisco. For most update types, you can schedule their download and installation; see Scheduling Tasks, page 42-1 and Using Recurring Rule Updates, page 46-13.

Table 46-1 ASA FirePOWER Module Update Types

Update Type	Description	Schedule?	Uninstall?
patches	Patches include a limited range of fixes (and usually change the fourth digit in the version number; for example, 5.4.0.1).	yes	yes
feature updates	Feature updates are more comprehensive than patches and generally include new features (and usually change the third digit in the version number; for example, 5.4.1).	yes	yes
major updates (major and minor version releases)	Major updates, sometimes referred to as upgrades, include new features and functionality and may entail large-scale changes (and usually change the first or second digit in the version number; for example, 5.3 or 5.4).	no	no
VDB	VDB updates affect the database of known vulnerabilities to which hosts may be susceptible.	yes	no
intrusion rules	Intrusion rule updates provide new and updated intrusion rules and preprocessor rules, modified states for existing rules, and modified default intrusion policy settings. Rule updates may also delete rules, provide new rule categories and default variables, and modify default variable values.	yes	no
geolocation database (GeoDB)	GeoDB updates provide updated information on physical locations, connection types, and so on that your system can associate with detected routable IP addresses. You can use geolocation data as a condition in access control rules. You must install the GeoDB to view geolocation details.	yes	no

Note that while you can uninstall patches and other minor updates, you cannot uninstall major updates or return to previous versions of the VDB, GeoDB, or intrusion rules. If you updated to a new major version and you need to revert to an older version, contact Support.

Performing Software Updates

License: Any

There are a few basic steps to updating. First, you **must** prepare for the update by reading the release notes and completing any required pre-update tasks. Then, you can begin the update. You must verify the update's success. Finally, complete any required post-update steps.

For more information, see the following sections:

- Planning for the Update, page 46-2
- Understanding the Update Process, page 46-3
- Updating the ASA FirePOWER Module Software, page 46-4
- Monitoring the Status of Major Updates, page 46-6

Planning for the Update

License: Any

Before you begin the update, you must thoroughly read and understand the release notes, which you can download from the Support Site. The release notes describe new features and functionality, and known and resolved issues. The release notes also contain important information on prerequisites, warnings, and specific installation and uninstallation instructions.

The following sections provide an overview of some of the factors you must consider when planning for the update.

Software Version Requirements

You must make sure you are running the correct software version. The release notes indicate the required version. If you are running an earlier version, you can obtain updates from the Support Site.

Time and Disk Space Requirements

Make sure you have enough free disk space and allow enough time for the update. The release notes indicate space and time requirements.

Configuration Backup Guidelines

Before you begin a major update, Cisco recommends that you delete any backups that reside on the ASA FirePOWER module after copying them to an external location. Regardless of the update type, you should also back up current configuration data to an external location. See Using Backup and Restore, page 48-1.

When to Perform the Update

Because the update process may affect traffic inspection and traffic flow, and because the Data Correlator is disabled while an update is in progress, Cisco recommends you perform the update in a maintenance window or at a time when the interruption will have the least impact.

Understanding the Update Process

License: Any

You use the ASA FirePOWER module interface to update the ASA FirePOWER module.

The Product Updates page (**Configuration > ASA FirePOWER Configuration > Updates**) shows the version of each update, as well as the date and time it was generated. It also indicates whether a software reboot is required as part of the update. When you upload updates obtained from Support, they appear on the page. Uninstallers for patch and feature updates also appear; see Uninstalling Software Updates, page 46-7. The page can also list VDB updates.



Tip For patches and feature updates, you can take advantage of the automated update feature; see Automating Software Updates, page 42-5.

Traffic Flow and Inspection

When you install or uninstall updates, the following capabilities may be affected:

- traffic inspection, including application and user awareness and control, URL filtering, Security Intelligence filtering, intrusion detection and prevention, and connection logging
- traffic flow

The Data Correlator does not run during system updates. It resumes when the update is complete.

The manner and duration of network traffic interruption depends on how yourASA FirePOWER module is configured and deployed, and whether the update reboots the ASA FirePOWER module. For specific information on how and when network traffic is affected for a particular update, see the release notes.

Using the ASA FirePOWER Module During the Update

Regardless of the type of update, do **not** use the ASA FirePOWER module to perform tasks other than monitoring the update.

To prevent you from using the ASA FirePOWER module during a major update, and to allow you to easily monitor a major update's progress, the system streamlines the ASA FirePOWER module interface. You can monitor a minor update's progress in the task queue (**Monitoring > ASA FirePOWER Monitoring > Task Status**). Although you are not prohibited from using the ASA FirePOWER module during a minor update, Cisco recommends against it.

Even for minor updates, the ASA FirePOWER module may become unavailable during the update process. This is expected behavior. If this occurs, wait until you can again access the ASA FirePOWER module. If the update is still running, you **must** continue to refrain from using the ASA FirePOWER module until the update has completed. Note that while updating, the ASA FirePOWER module may reboot a second time; this is also expected behavior.

Caution If you encounter issues with the update (for example, if the update has failed or if a manual refresh of the Update Status page shows no progress), do **not** restart the update. Instead, contact Support.

After the Update

You **must** complete all of the post-update tasks listed in the release notes to ensure that your deployment is performing properly.

The most important post-update task is to reapply access control policies. Note that applying an access control policy may cause a short pause in traffic flow and processing, and may also cause a few packets to pass uninspected; see Deploying Configuration Changes, page 4-11.

Additionally, you should:

- verify that the update succeeded
- update your intrusion rules, VDB, and GeoDB, if necessary
- make any required configuration changes, based on the information in the release notes
- perform any additional post-update tasks listed in the release notes

Updating the ASA FirePOWER Module Software

License: Any

Update the ASA FirePOWER module software in one of two ways, depending on the type of update and whether your ASA FirePOWER module has access to the Internet:

- You can obtain the update directly from the Support Site if your ASA FirePOWER module has access to the Internet. This option is **not** supported for major updates.
- You can manually download the update from the Support Site and then upload it to the ASA FirePOWER module. Choose this option if your ASA FirePOWER module does not have access to the Internet or if you are performing a major update.

For major updates, updating the ASA FirePOWER module removes uninstallers for previous updates.

To update the ASA FirePOWER Module Software:

Step 1 Read the release notes and complete any required pre-update tasks.

Pre-update tasks may include making sure that: the ASA FirePOWER module is running the correct version of the Cisco software, you have enough free disk space to perform the update, you set aside adequate time to perform the update, you backed up configuration data, and so on.

Step 2 Upload the update. You have two options, depending on the type of update and whether your ASA FirePOWER module has access to the Internet:

- For all except major updates, and if your ASA FirePOWER module has access to the Internet, select **Configuration > ASA FirePOWER Configuration > Updates**, then click **Download Updates** to check for the latest updates on either of the following Support Sites:
  - **Sourcefire:** (https://support.sourcefire.com/)
  - **Cisco:** (http://www.cisco.com/cisco/web/support/index.html)
- For major updates, or if your ASA FirePOWER module does not have access to the Internet, you must first manually download the update from either of the following Support Sites:
  - **Sourcefire:** (https://support.sourcefire.com/)
  - **Cisco:** (http://www.cisco.com/cisco/web/support/index.html)
- Select **Configuration > ASA FirePOWER Configuration > Updates**, then click **Upload Update**. Click **Choose File** to navigate to and select the update and click **Upload**.

Note Download the update directly from the Support Site, either manually or by clicking **Download Updates** on the Product Updates tab. If you transfer an update file by email, it may become corrupted.

The update is uploaded.

Step 3 Select **Monitoring > ASA FirePOWER Monitoring > Task Status** to view the task queue and make sure that there are no jobs in process.

Tasks that are running when the update begins are stopped and cannot be resumed; you must manually delete them from the task queue after the update completes. The task queue automatically refreshes every 10 seconds. You must wait until any long-running tasks are complete before you begin the update.

Step 4 Select **Configuration > ASA FirePOWER Configuration > Updates**.

The Product Updates page appears.

Step 5 Click the install icon next to the update you uploaded.

The update process begins. How you monitor the update depends on whether the update is a major or minor update. See the ASA FirePOWER Module Update Types table and the release notes to determine your update type:

- For minor updates, you can monitor the update's progress in the task queue (**Monitoring > ASA FirePOWER Monitoring > Task Status**).
- For major updates, you can begin monitoring the update's progress in the task queue. However, after the ASA FirePOWER module completes its necessary pre-update checks, you are locked out of the module interface. When you regain access, the Upgrade Status page appears. See Monitoring the Status of Major Updates, page 46-6 for information.

Caution Regardless of the update type, do **not** use the ASA FirePOWER module to perform tasks other than monitoring the update until the update has completed and, if necessary, the ASA FirePOWER module reboots. For more information, see Using the ASA FirePOWER Module During the Update, page 46-4.

Step 6 After the update finishes, access the ASA FirePOWER module interface and refresh the page. Otherwise, the interface may exhibit unexpected behavior. If you are the first user to access the interface after a major update, the End User License Agreement (EULA) may appear. You must review and accept the EULA to continue.

Step 7 If the rule update available on the Support Site is newer than the rules on your ASA FirePOWER module, import the newer rules.

For more information, see Importing Rule Updates and Local Rule Files, page 46-9.

Step 8 Reapply access control policies.

Applying an access control policy may cause a short pause in traffic flow and processing, and may also cause a few packets to pass uninspected. For more information, see Deploying Configuration Changes, page 4-11.

Step 9 If the VDB available on the Support Site is newer than the most recently installed VDB, install the latest VDB.

Installing a VDB update causes a short pause in traffic flow and processing, and may also cause a few packets to pass uninspected. For more information, see Updating the Vulnerability Database, page 46-8.

Monitoring the Status of Major Updates

License: Any

For major updates, the ASA FirePOWER module provides you with a streamlined interface so that you can easily monitor the update process. The streamlined interface also prevents you from using the ASA FirePOWER module to perform tasks other than monitoring the update.You can begin monitoring the update's progress in the task queue (**Monitoring > ASA FirePOWER Monitoring > Task Status**). However, after the ASA FirePOWER module completes its necessary pre-update checks, you are locked out of the user interface until a streamlined update page appears.

The streamlined interface displays the version you are updating from, the version you are updating to, and the time that has elapsed since the update began. It also displays a progress bar and gives details about the script currently running.

Tip Click **show log for current script** to see the update log. Click **hide log for current script** to hide the log again.

If the update fails for any reason, the page displays an error message indicating the time and date of the failure, which script was running when the update failed, and instructions on how to contact Support. Do **not** restart the update.

Caution If you encounter any other issue with the update (for example, if a manual refresh of the page shows no progress for an extended period of time), do **not** restart the update. Instead, contact Support.

When the update completes, the ASA FirePOWER module displays a success message and reboots. After the ASA FirePOWER module finishes rebooting, complete any required post-update steps.

Uninstalling Software Updates

License: Any

When you apply a patch or feature update, the update process creates an uninstaller that allows you to remove the update.

When you uninstall an update, the resulting Cisco software version depends on the update path. For example, consider a scenario where you updated directly from Version 5.0 to Version 5.0.0.2. Uninstalling the Version 5.0.0.2 patch might result in Version 5.0.0.1, even though you never installed the Version 5.0.0.1 update. For information on the resulting Cisco software version when you uninstall an update, see the release notes.

Note Uninstalling is not supported for major updates. If you updated to a new major version and you need to revert to an older version, contact Support.

Traffic Flow and Inspection

Uninstalling an update may affect traffic inspection and traffic flow. For specific information on how and when network traffic is affected for a particular update, see the release notes.

After the Uninstallation

After you uninstall the update, verify that the uninstall succeeded. For specific information for each update, see the release notes.

To uninstall a patch or feature update:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Updates**.

The Product Updates page appears.

Step 2 Click the install icon next to the uninstaller for the update you want to remove.

If prompted, confirm that you want to uninstall the update and reboot the ASA FirePOWER module.

The uninstall process begins. You can monitor its progress in the task queue (**Monitoring > ASA FirePOWER Monitoring > Task Status**).

Caution Do **not** use the ASA FirePOWER module interface to perform tasks until the uninstall has completed and, if necessary, the ASA FirePOWER module reboots. For more information, see Using the ASA FirePOWER Module During the Update, page 46-4.

Step 3 Refresh the page. Otherwise, the interface may exhibit unexpected behavior.

Updating the Vulnerability Database

License: Any

The Cisco Vulnerability Database (VDB) is a database of known vulnerabilities to which hosts may be susceptible. The Cisco Vulnerability Research Team (VRT) issues periodic updates to the VDB. To update the VDB, use the Product Updates page.

Note Installing a VDB update with detection updates may cause a short pause in traffic flow and processing, and may also cause a few packets to pass uninspected.You may want to schedule the update during low system usage times to minimize the impact of any system downtime.

Note After you complete a VDB update, reapply any out-of-date access control policy. Keep in mind that installing a VDB or reapplying an access control policy can cause a short pause in traffic flow and processing, and may also cause a few packets to pass uninspected. For more information, see Deploying Configuration Changes, page 4-11.

This section explains how to plan for and perform manual VDB updates.

To update the vulnerability database:

Step 1 Read the VDB Update Advisory Text for the update.

The advisory text includes information about the changes to the VDB made in the update.

Step 2 Select **Configuration > ASA FirePOWER Configuration > Updates**.

The Product Updates page appears.

Step 3 Upload the update:

- If your ASA FirePOWER module has access to the Internet, click **Download Updates** to check for the latest updates on either of the following Support Sites:
  - **Sourcefire:** (https://support.sourcefire.com/)
  - **Cisco:** (http://www.cisco.com/cisco/web/support/index.html)
- If your ASA FirePOWER module does not have access to the Internet, manually download the update from one of the following Support Sites, then click **Upload Update**. Click **Choose File** to navigate to and select the update and click **Upload**:
  - **Sourcefire:** (https://support.sourcefire.com/)
  - **Cisco:** (http://www.cisco.com/cisco/web/support/index.html)

Note Download the update directly from the Support Site either manually or by clicking **Download Updates**. If you transfer an update file by email, it may become corrupted.

The update is uploaded.

Step 4 Click the install icon next to the VDB update.

The Install Update page appears.

Step 5 Click **Install**.

The update process begins. You can monitor the update's progress in the task queue (**Monitoring > ASA FirePOWER Monitoring > Task Status**).

Caution If you encounter issues with the update (for example, if the task queue indicates that the update has failed) **do not** restart the update. Instead, contact Support.

You must reapply any out-of-date access control policies for the VDB update to take effect; see Deploying Configuration Changes, page 4-11.

Importing Rule Updates and Local Rule Files

License: Any

As new vulnerabilities become known, the Cisco Vulnerability Research Team (VRT) releases rule updates that you can first import onto your ASA FirePOWER module, then implement by applying affected access control, network analysis, and intrusion policies.

Rule updates are cumulative, and Cisco recommends you always import the latest update. You cannot import a rule update that either matches or predates the version of the currently installed rules.

Note Rule updates may contain new binaries, so make sure your process for downloading and installing them complies with your security policies. In addition, rule updates may be large, so import rules during periods of low network use.

A rule update may provide the following:

- **new and modified rules and rule states**—Rule updates provide new and updated intrusion and preprocessor rules. For new rules, the rule state may be different in each system-provided intrusion policy. For example, a new rule may be enabled in the Security over Connectivity intrusion policy and disabled in the Connectivity over Security intrusion policy. Rule updates may also change the default state of existing rules, or delete existing rules entirely.
- **new rule categories**—Rule updates may include new rule categories, which are always added.
- **modified preprocessor and advanced settings**—Rule updates may change the advanced settings in the system-provided intrusion policies and the preprocessor settings in system-provided network analysis policies. They can also update default values for the advanced preprocessing and performance options in your access control policies.
- **new and modified variables**—Rule updates may modify default values for existing default variables, but do not override your changes. New variables are always added.

Understanding When Rule Updates Modify Policies

Rule updates can affect both system-provided and custom network analysis policies, as well as all access control policies:

- **system provided**—Changes to system-provided network analysis and intrusion policies, as well as any changes to advanced access control settings, automatically take effect when you reapply the policies after the update.

- **custom**—Because every custom network analysis and intrusion policy uses a system-provided policy as its base, or as the eventual base in a policy chain, rule updates can affect custom network analysis and intrusion policies. However, you can prevent rule updates from automatically making those changes. This allows you to update system-provided base policies manually, on a schedule independent of rule update imports. Regardless of your choice (implemented on a per-custom-policy basis), updates to system-provided policies do **not** override any settings you customized. For more information, see Allowing Rule Updates to Modify a System-Provided Base Policy, page 19-4.

Note that importing a rule update discards all cached changes to network analysis and intrusion policies. For your convenience, the Rule Updates page lists policies with cached changes. For more information, see Resolving Conflicts and Committing Policy Changes, page 18-15.

Reapplying Policies

For changes made by a rule update to take affect, you must reapply any modified policies. When importing a rule update, you can configure the system to automatically reapply intrusion or access control policies. This is especially useful if you allow the rule update to modify system-provided base policies.

- Reapplying an access control policy also reapplies associated SSL, network analysis, and file policies, but does **not** reapply intrusion policies. It also updates the default values for any modified advanced settings. Because you cannot apply a network analysis policy independently, you **must** reapply access control policies if you want to update preprocessor settings in network analysis policies.
- Reapplying intrusion policies allows you to update rules and other changed intrusion policy settings. You can reapply intrusion policies in conjunction with access control policies, or you can apply only intrusion policies to update intrusion rules without updating any other access control configurations.

When a rule update includes shared object rules, applying an access control or intrusion policy for the first time after the import causes a short pause in traffic flow and processing, and may also cause a few packets to pass uninspected. For more information on applying access control and intrusion policies, including requirements, other effects, and recommendations, see Deploying Configuration Changes, page 4-11.

For more information on importing rule updates, see:

- Using One-Time Rule Updates, page 46-10 explains how to import a single rule update from the Support Site.
- Using Recurring Rule Updates, page 46-13 explains how to use an automated feature to download and install rule updates from the Support Site.
- Importing Local Rule Files, page 46-14 explains how to import a copy of a standard text rules file that you have created on a local machine.
- Viewing the Rule Update Log, page 46-15 explains the rule update log.

Using One-Time Rule Updates

License: Any

There are two methods that you can use for one-time rule updates:

- Using Manual One-Time Rule Updates, page 46-11 explains how to manually download a rule update from the Support Site and then manually install the rule update.
- Using Automatic One-Time Rule Updates, page 46-12 explains how to use an automated feature to search the Support Site for new rule updates and upload them.

Using Manual One-Time Rule Updates

License: Any

The following procedure explains how to import a new rule update manually. This procedure is especially useful if your ASA FirePOWER module does not have Internet access.

To manually import a rule update:

Step 1 From a computer that can access the Internet, access either of the following sites:

- **Sourcefire:** (https://support.sourcefire.com/)
- **Cisco:** (http://www.cisco.com/cisco/web/support/index.html)

Step 2 Click **Download**, then click **Rules**.

Step 3 Navigate to the latest rule update.

Rule updates are cumulative; you cannot import a rule update that either matches or predates the version of the currently installed rules.

Step 4 Click the rule update file that you want to download and save it to your computer.

Step 5 Select **Configuration > ASA FirePOWER Configuration > Updates**, then select the **Rule Updates** tab.

The Rule Updates page appears.

Tip You can also click **Import Rules** on the Rule Editor page **(Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy > Rule Editor)**.

Step 6 Optionally, click **Delete All Local Rules**, then click **OK** to move all user-defined rules that you have created or imported to the deleted folder. See Deleting Custom Rules, page 30-105 for more information.

Step 7 Select **Rule Update or text rule file to upload and install** and click **Choose File** to navigate to and select the rule update file.

Step 8 Optionally, reapply policies after the update completes:

- Select **Reapply intrusion policies after the rule update import completes** to automatically reapply intrusion policies. Choose only this option to update rules and other changed intrusion policy settings without having to update any other access control configurations you may have made. You **must** select this option to reapply intrusion policies in conjunction with access control policies; reapplying access control policies in this case does not perform a complete apply.
- Select **Reapply access control policies after the rule update import completes** to automatically reapply access control policies and their associated SSL, network analysis, and file policies, but not intrusion policies. Selecting this option also updates the default values for any modified access control advanced settings. Because you cannot apply a network analysis policy independently of its parent access control policy, you **must** reapply access control policies if you want to update preprocessor settings in network analysis policies.

Step 9 Click **Import**.

The system installs the rule update and displays the Rule Update Log detailed view; see Understanding the Rule Update Import Log Detailed View, page 46-18. The system also applies policies as you specified in the previous step; see Deploying Configuration Changes, page 4-11 and Applying an Intrusion Policy, page 26-7.

Note Contact Support if you receive an error message while installing the rule update.

Using Automatic One-Time Rule Updates

License: Any

The following procedure explains how to import a new rule update by automatically connecting to the Support Site. You can use this procedure only if the ASA FirePOWER module has Internet access.

To automatically import a rule update:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Updates**, then select the **Rule Updates** tab.

The Rule Updates page appears.

Tip You can also click **Import Rules** on the Rule Editor page (**Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy > Rule Editor)**.

Step 2 Optionally, click **Delete All Local Rules**, then click **OK** to move all user-defined rules that you have created or imported to the deleted folder. See Deleting Custom Rules, page 30-105 for more information.

Step 3 Select **Download new Rule Update from the Support Site**.

Step 4 Optionally, reapply policies after the update completes:

- Select **Reapply intrusion policies after the rule update import completes** to automatically reapply intrusion policies. Choose only this option to update rules and other changed intrusion policy settings without having to update any other access control configurations you may have made. You **must** select this option to reapply intrusion policies in conjunction with access control policies; reapplying access control policies in this case does not perform a complete apply.
- Select **Reapply access control policies after the rule update import completes** to automatically reapply access control, network analysis, and file policies, but not intrusion policies. Selecting this option also updates the default values for any modified access control advanced settings. Because you cannot apply a network analysis policy independently of its parent access control policy, you **must** reapply access control policies if you want to update preprocessor settings in network analysis policies.

Step 5 Click **Import**.

The system installs the rule update and displays the Rule Update Log detailed view; see Understanding the Rule Update Import Log Detailed View, page 46-18. The system also applies policies as you specified in the previous step; see Deploying Configuration Changes, page 4-11 and Applying an Intrusion Policy, page 26-7.

Note Contact Support if you receive an error message while installing the rule update.

Using Recurring Rule Updates

License: Any

You can import rule updates on a daily, weekly, or monthly basis, using the Rule Updates page.

Applicable subtasks in the rule update import occur in the following order: download, install, base policy update, and policy reapply. When one subtask completes, the next subtask begins. Note that you can only apply policies previously applied by the ASA FirePOWER module where the recurring import is configured.

To schedule recurring rule updates:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Updates**, then select the **Rule Updates** tab.

The Rule Updates page appears.

Tip You can also click **Import Rules** on the Rule Editor page (**Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy > Rule Editor**).

Step 2 Optionally, click **Delete All Local Rules**, then click **OK** to move all user-defined rules that you have created or imported to the deleted folder. See Deleting Custom Rules, page 30-105 for more information.

Step 3 Select **Enable Recurring Rule Update Imports**.

The page expands to display options for configuring recurring imports. Import status messages appear beneath the **Recurring Rule Update Imports** section heading. Recurring imports are enabled when you save your settings.

Tip To disable recurring imports, clear the **Enable Recurring Rule Update Imports** check box and click **Save**.

Step 4 In the **Import Frequency** field, select **Daily**, **Weekly**, or **Monthly** from the drop-down list.

If you selected a weekly or monthly import frequency, use the drop-down lists that appear to select the day of the week or month when you want to import rule updates. Select from a recurring task drop-down list either by clicking or by typing the first letter or number of your selection one or more times and pressing Enter.

Step 5 In the **Import Frequency** field, specify the time when you want to start your recurring rule update import.

Step 6 Optionally, reapply policies after the update completes:

- Select **Reapply intrusion policies after the rule update import completes** to automatically reapply intrusion policies. Choose only this option to update rules and other changed intrusion policy settings without having to update any other access control configurations you may have made. You **must** select this option to reapply intrusion policies in conjunction with access control policies; reapplying access control policies in this case does not perform a complete apply.
- Select **Reapply access control policies after the rule update import completes** to automatically reapply access control policies and their associated SSL, network analysis, and file policies, but not intrusion policies. Selecting this option also updates the default values for any modified access control advanced settings. Because you cannot apply a network analysis policy independently of its parent access control policy, you **must** reapply access control policies if you want to update preprocessor settings in network analysis policies.

Step 7 Click **Save** to enable recurring rule update imports using your settings.

The status message under the Recurring Rule Update Imports section heading changes to indicate that the rule update has not yet run. At the scheduled time, the system installs the rule update and applies policies as you specified in the previous step; see Deploying Configuration Changes, page 4-11 and Applying an Intrusion Policy, page 26-7.

You can log off or perform other tasks before or during the import. When accessed during an import, the Rule Update Log displays a red status icon ( ), and you can view messages as they occur in the Rule Update Log detailed view. Depending on the rule update size and content, several minutes may pass before status messages appear. For more information, see Viewing the Rule Update Log, page 46-15.

Note Contact Support if you receive an error message while installing the rule update.

Importing Local Rule Files

License: Any

A local rule is a custom standard text rule that you import from a local machine as a plain text file with ASCII or UTF-8 encoding. You can create local rules using the instructions in the Snort users manual, which is available at http://www.snort.org.

Note the following regarding importing local rules:

- The text file name can include alphanumeric characters, spaces, and no special characters other than underscore (_), period (.), and dash (-).
- You do not have to specify a Generator ID (GID); if you do, you can specify only GID 1 for a standard text rule or 138 for a sensitive data rule.
- Do **not** specify a Snort ID (SID) or revision number when importing a rule for the first time; this avoids collisions with SIDs of other rules, including deleted rules.

  The system will automatically assign the rule the next available custom rule SID of 1000000 or greater, and a revision number of 1.
- You **must** include the SID assigned by the system and a revision number greater than the current revision number when importing an updated version of a local rule that you have previously imported.

  To view the revision number for a current local rule, display the Rule Editor page (**Policies > Intrusion Policy > Rule Editor**), click on the local rule category to expand the folder, then click **Edit** next to the rule.
- You can reinstate a local rule that you have deleted by importing the rule using the SID assigned by the system and a revision number greater than the current revision number. Note that the system automatically increments the revision number when you delete a local rule; this is a device that allows you to reinstate local rules.

  To view the revision number for a deleted local rule, display the Rule Editor page (**Policies > Intrusion Policy > Rule Editor**), click on the deleted rule category to expand the folder, then click **Edit** next to the rule.
- You cannot import a rule file that includes a rule with a SID greater than 2147483647; the import will fail.
- If you import a rule that includes a list of source or destination ports that is longer than 64 characters, the import will fail.

- The system always sets local rules that you import to the disabled rule state; you must manually set the state of local rules before you can use them in your intrusion policy. See Setting Rule States, page 27-19 for more information.
- You must make sure that the rules in the file do not contain any escape characters.
- The rules importer requires that all custom rules are imported in ASCII or UTF-8 encoding.
- All imported local rules are automatically saved in the local rule category.
- All deleted local rules are moved from the local rule category to the deleted rule category.
- The system imports local rules preceded with a single pound character (#).
- The system ignores local rules preceded with two pound characters (##) and does not import them.
- Policy validation fails if you enable an imported local rule that uses the deprecated `threshold` keyword in combination with the intrusion event thresholding feature in an intrusion policy. See Configuring Event Thresholding, page 27-21 for more information.

To import local rule files:

Step 1 Select **Policies > Intrusion Policy > Rule Editor**.

The Rule Editor page appears.

Step 2 Click **Import Rules**.

The Import Rules page appears.

Tip You can also select **System > Updates**, then select the **Rule Updates** tab.

Step 3 Select **Rule Update or text rule file to upload and install** and click **Choose File** to navigate to the rule file. Note that all rules uploaded in this manner are saved in the local rule category.

Tip You can import **only** plain text files with ASCII or UTF-8 encoding.

Step 4 Click **Import**.

The rule file is imported. Make sure you enable the appropriate rules in your intrusion policies. The rules are not activated until the next time you apply the affected policies.

Note The system does **not** use the new rule set for inspection until after you apply your intrusion policies. See Deploying Configuration Changes, page 4-11 for procedures.

Viewing the Rule Update Log

License: Any

The ASA FirePOWER module generates a record for each rule update and local rule file that you import.

Each record includes a time stamp, the name of the user who imported the file, and a status icon indicating whether the import succeeded or failed. You can maintain a list of all rule updates and local rule files that you import, delete any record from the list, and access detailed records for all imported rules and rule update components. Actions you can take in the Rule Update Log are described in the following table.

Table 46-2 Rule Update Log Actions

To...	You can...
learn more about the contents of the columns in the table	find more information in Understanding the Rule Update Log Table, page 46-16.
delete an import file record from the import log, including detailed records for all objects included with the file	click the delete icon ( ) next to the file name for the import file. **Note** Deleting the file from the log does not delete any object imported in the import file, but only deletes the import log records.
view details for each object imported in a rule update or local rule file	click the view icon ( ) next to the file name for the import file.

See the following sections for more information:

- Understanding the Rule Update Log Table, page 46-16 describes the fields in the list of rule updates and local rule files that you import.
- Viewing Rule Update Import Log Details, page 46-17 describes the detailed record for each object imported in a rule update or local rule file.
- Understanding the Rule Update Import Log Detailed View, page 46-18 describes each field in the Rule Update Log detailed view.

To view the Rule Update Log:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Updates**, then select the **Rule Updates** tab.

The Rule Updates page appears.

Tip You can also click **Import Rules** on the Rule Editor page (**Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy > Rule Editor**).

Step 2 Click **Rule Update Log**.

The Rule Update Log page appears. This page lists each imported rule update and local rule file.

Understanding the Rule Update Log Table

License: Any

The fields in the list of rule updates and local rule files that you import are described in the following table.

Table 46-3 Rule Update Log Fields

Field	Description
Summary	The name of the import file. If the import fails, a brief statement of the reason for the failure appears under the file name.
Time	The time and date that the import started.
User ID	The user name of the user that triggered the import.
Status	Whether the import: • succeeded ( ) • failed or is currently in progress ( ) **Tip** The red status icon indicating an unsuccessful or incomplete import appears on the Rule Update Log page during the import and is replaced by the green icon only when the import has successfully completed.

Click the view icon ( ) next to the rule update or file name to view the Rule Update Log detailed page for the rule update or local rule file, or click the delete icon ( ) to delete the file record and all detailed object records imported with the file.

Tip You can view import details as they appear while a rule update import is in progress.

Viewing Rule Update Import Log Details

License: Any

The Rule Update Import Log detailed view lists a detailed record for each object imported in a rule update or local rule file. You can also create a custom workflow or report from the records listed that includes only the information that matches your specific needs.

The following table describes specific actions you can perform on a Rule Update Import Log detailed view.

Table 46-4 Rule Update Import Log Detailed View Actions

To...	You can...
learn more about the contents of the columns in the table	find more information in Understanding the Rule Update Import Log Detailed View, page 46-18.

To view the Rule Update Import Log Detailed View:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Updates**, then select the **Rule Updates** tab.

The Rule Updates page appears.

Tip You can also click **Import Rules** on the Rule Editor page (**Configuration > ASA FirePOWER Configuration > Policies > Intrusion Policy > Rule Editor**).

Step 2 Click **Rule Update Log**.

The Rule Update Log page appears.

Step 3 Click the view icon ( ) next to the file whose detailed records you want to view.

The table view of detailed records appears.

Understanding the Rule Update Import Log Detailed View

License: Any

You can view a detailed record for each object imported in a rule update or local rule file. The fields in the Rule Update Log detailed view are described in the following table.

Table 46-5 Rule Update Import Log Detailed View Fields

Field	Description
Time	The time and date the import began.
Name	The name of the imported object, which for rules corresponds to the rule Message field, and for rule update components is the component name.
Type	The type of imported object, which can be one of the following: • `rule update component` (an imported component such as a rule pack or policy pack) • `rule` (for rules, a new or updated rule; note that in Version 5.0.1 this value replaced the `update` value, which is deprecated) • `policy apply` (the **Reapply intrusion policies after the Rule Update import completes** option was enabled for the import)
Action	An indication that one of the following has occurred for the object type: • `new` (for a rule, this is the first time the rule has been stored on this ASA FirePOWER module) • `changed` (for a rule update component or rule, the rule update component has been modified, or the rule has a higher revision number and the same GID and SID) • `collision` (for a rule update component or rule, import was skipped because its revision conflicts with an existing component or rule) • `deleted` (for rules, the rule has been deleted from the rule update) • `enabled` (for a rule update edit, a preprocessor, rule, or other feature has been enabled in a system-provided policy) • `disabled` (for rules, the rule has been disabled in a system-provided policy) • `drop` (for rules, the rule has been set to Drop and Generate Events in a system-provided policy) • `error` (for a rule update or local rule file, the import failed) • `apply` (the **Reapply intrusion policies after the Rule Update import completes** option was enabled for the import)
Default Action	The default action defined by the rule update. When the imported object type is `rule`, the default action is `Pass`, `Alert`, or `Drop`. For all other imported object types, there is no default action.
GID	The generator ID for a rule. For example, `1` (standard text rule) or `3` (shared object rule).
SID	The SID for a rule.
Rev	The revision number for a rule.

Table 46-5 Rule Update Import Log Detailed View Fields (continued)

Field	Description
Policy	For imported rules, this field displays `All`, which indicates that the imported rule was included in all system-provided intrusion policies. For other types of imported objects, this field is blank.
Details	A string unique to the component or rule. For rules, the GID, SID, and previous revision number for a changed rule, displayed as `previously (GID:SID:Rev)`. This field is blank for a rule that has not changed.
Count	The count (`1`) for each record. The Count field appears in a table view when the table is constrained, and the Rule Update Log detailed view is constrained by default to rule update records.

Updating the Geolocation Database

License: Any

The Cisco Geolocation Database (GeoDB) is a database of geographical data associated with routable IP addresses. The ASA FirePOWER module provides the country and continent. When your system detects GeoDB information that matches a detected IP address, you can view the geolocation information associated with that IP address. Cisco issues periodic updates to the GeoDB.

To update the GeoDB, use the Geolocation Updates page (**Configuration > ASA FirePOWER Configuration > Updates > Geolocation Updates**). When you upload GeoDB updates, they appear on this page.

The installation usually takes 30 to 40 minutes. Although a GeoDB update does not interrupt any other system functions (including the ongoing collection of geolocation information), the update does consume system resources while it completes. Consider this when planning your updates.

This section explains how to plan for and perform manual GeoDB updates. You can also take advantage of the automated update feature to schedule GeoDB updates; for more information, see Automating Geolocation Database Updates, page 42-4.

To update the geolocation database:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Updates**.

The Product Updates page appears.

Step 2 Click the **Geolocation Updates** tab.

The Geolocation Updates page appears.

Step 3 Upload the update.

- If your ASA FirePOWER module has access to the Internet, click **Download and install geolocation update from the Support Site** to check for the latest updates on either of the following Support Sites:
  - **Sourcefire:** (https://support.sourcefire.com/)
  - **Cisco:** (http://www.cisco.com/cisco/web/support/index.html)
- If your ASA FirePOWER module does not have access to the Internet, manually download the update from either of the Support Sites, then click **Upload and install geolocation update**. Click **Choose File** to navigate to and select the update and click **Import**:
  - **Sourcefire:** (https://support.sourcefire.com/)
  - **Cisco:** (http://www.cisco.com/cisco/web/support/index.html)



Note Download the update directly from the Support Site, either manually or by clicking **Download and install geolocation update from the Support Site** on the Geolocation Updates page. If you transfer an update file by email, it may become corrupted.

The update process begins. The average duration of update installation is 30 to 40 minutes. You can monitor the update's progress in the task queue (**Monitoring > ASA FirePOWER Monitoring > Task Status**).

Step 4 After the update finishes, return to the Geolocation Updates page to confirm that the GeoDB build number matches the update you installed.

The GeoDB update overrides any previous versions of the GeoDB and is effective immediately. Although it may take a few minutes for a GeoDB update to take effect throughout your deployment, you do not have to reapply access control policies after you update.



Monitoring the System

The ASA FirePOWER module provides many useful monitoring features to assist you in the daily administration of your system, all on a single page. For example, on the Host Statistics page you can monitor basic host statistics. The following sections provide more information about the monitoring features that the system provides:

- Viewing Host Statistics, page 47-1 describes how to view host information such as:
- system uptime
- disk and memory usage
- system processes
- intrusion event information
- Monitoring System Status and Disk Space Usage, page 47-2 describes how to view basic event and disk partition information.
- Viewing System Process Status, page 47-2 describes how to view basic process status.
- Understanding Running Processes, page 47-4 describes the basic system processes that run on the appliance.

Viewing Host Statistics

License: Any

The Statistics page lists the current status of the following:

- general host statistics; see the Host Statistics table for details
- intrusion event information (requires Protection); see Viewing Events, page 37-1 or details

The following table describes the host statistics listed on the Statistics page.

Table 47-1 Host Statistics

Category	Description
Time	The current time on the system.
Uptime	The number of days (if applicable), hours, and minutes since the system was last started.
Memory Usage	The percentage of system memory that is being used.

Table 47-1 Host Statistics (continued)

Category	Description
Load Average	The average number of processes in the CPU queue for the past 1 minute, 5 minutes, and 15 minutes.
Disk Usage	The percentage of the disk that is being used. Click the arrow to view more detailed host statistics. See Monitoring System Status and Disk Space Usage, page 47-2 for more information.
Processes	A summary of the processes running on the system. See Viewing System Process Status, page 47-2 for more information.

To view the Statistics page:

Step 1 Select **Monitoring > ASA FirePOWER Monitoring > Statistics**.

The Statistics page appears.

Monitoring System Status and Disk Space Usage

License: Any

The Disk Usage section of the Statistics page provides a quick synopsis of disk usage, both by category and by partition status. If you have a malware storage pack installed on a device, you can also check its partition status. You can monitor this page from time to time to ensure that enough disk space is available for system processes and the database.

To access disk usage information:

Step 1 Select **Monitoring > ASA FirePOWER Monitoring > Statistics**.

The Statistics page appears.

For more information on the disk usage categories, see Understanding the Disk Usage Widget, page 40-3.

Step 2 Click the down arrow next to **Total** to expand it.

The Disk Usage section expands, displaying partition usage. If you have a malware storage pack installed, the `/var/storage` partition usage is also displayed.

Viewing System Process Status

License: Any

The Processes section of the Host Statistics page allows you to see the processes that are currently running on an appliance. It provides general process information and specific information for each running process.

The following table describes each column that appears in the process list.

Table 47-2 Process Status

Column	Description
Pid	The process ID number
Username	The name of the user or group running the process
Pri	The process priority
Nice	The *nice* value, which is a value that indicates the scheduling priority of a process. Values range between -20 (highest priority) and 19 (lowest priority)
Size	The memory size used by the process (in kilobytes unless the value is followed by `m`, which indicates megabytes)
Res	The amount of resident paging files in memory (in kilobytes unless the value is followed by `m`, which indicates megabytes)
State	The process state: • D — process is in uninterruptible sleep (usually Input/Output) • N — process has a positive nice value • R — process is runnable (on queue to run) • S — process is in sleep mode • T — process is being traced or stopped • W — process is paging • X — process is dead • Z — process is defunct • < — process has a negative nice value
Time	The amount of time (in hours:minutes:seconds) that the process has been running
Cpu	The percentage of CPU that the process is using
Command	The executable name of the process

To expand the process list:

Step 1 Select **Monitoring > ASA FirePOWER Monitoring > Statistics**.

The Statistics page appears.

Step 2 Click the down arrow next to **Processes**.

The process list expands, listing general process status information that includes the number and types of running tasks, the current time, the current system uptime, the system load average, CPU, memory, and swap information, and specific information about each running process.

Cpu(s) lists the following CPU usage information:

- user process usage percentage
- system process usage percentage
- nice usage percentage (CPU usage of processes that have a negative nice value, indicating a higher priority)

Nice values indicate the scheduled priority for system processes and can range between -20 (highest priority) and 19 (lowest priority).

- idle usage percentage

Mem lists the following memory usage information:

- total number of kilobytes in memory
- total number of used kilobytes in memory
- total number of free kilobytes in memory
- total number of buffered kilobytes in memory

Swap lists the following swap usage information:

- total number of kilobytes in swap
- total number of used kilobytes in swap
- total number of free kilobytes in swap
- total number of cached kilobytes in swap



Note For more information about the types of processes that run on the appliance, see Understanding Running Processes, page 47-4.

To collapse the process list:

Step 1 Click the up arrow next to **Processes**.

The process list collapses.

Understanding Running Processes

License: Any

There are two different types of processes that run on an appliance: daemons and executable files. Daemons always run, and executable files are run when required.

See the following sections for more information:

- Understanding System Daemons, page 47-4
- Understanding Executables and System Utilities, page 47-5

Understanding System Daemons

License: Any

Daemons continually run on an appliance. They ensure that services are available and spawn processes when required. The following table lists daemons that you may see on the Process Status page and provides a brief description of their functionality.

Note The table below is not an exhaustive list of all processes that may run on an appliance.

Table 47-3 System Daemons

Daemon	Description
crond	Manages the execution of scheduled commands (cron jobs)
dhclient	Manages dynamic host IP addressing
httpd	Manages the HTTP (Apache web server) process
httpsd	Manages the HTTPS (Apache web server with SSL) service, and checks for working SSL and valid certificate authentication; runs in the background to provide secure web access to the appliance
keventd	Manages Linux kernel event notification messages
klogd	Manages the interception and logging of Linux kernel messages
kswapd	Manages Linux kernel swap memory
kupdated	Manages the Linux kernel update process, which performs disk synchronization
mysqld	Manages ASA FirePOWER module database processes
ntpd	Manages the Network Time Protocol (NTP) process
pm	Manages all Cisco processes, starts required processes, restarts any process that fails unexpectedly
reportd	Manages reports
safe_mysqld	Manages safe mode operation of the database; restarts the database daemon if an error occurs and logs runtime information to a file
sfmgr	Provides the RPC service for remotely managing and configuring an appliance using an sftunnel connection to the appliance
sftroughd	Listens for connections on incoming sockets and then invokes the correct executable (typically the Cisco message broker, sfmb) to handle the request
sftunnel	Provides the secure communication channel for all processes requiring communication with a remote appliance
sshd	Manages the Secure Shell (SSH) process; runs in the background to provide SSH access to the appliance
syslogd	Manages the system logging (syslog) process

Understanding Executables and System Utilities

License: Any

There are a number of executables on the system that run when executed by other processes or through user action. The following table describes the executables that you may see on the Process Status page.

Table 47-4 System Executables and Utilities

Executable	Description
awk	Utility that executes programs written in the `awk` programming language
bash	GNU Bourne-Again SHell
cat	Utility that reads files and writes content to standard output
chown	Utility that changes user and group file permissions
chsh	Utility that changes the default login shell
cp	Utility that copies files
df	Utility that lists the amount of free space on the appliance
echo	Utility that writes content to standard output
egrep	Utility that searches files and folders for specified input; supports extended set of regular expressions not supported in standard grep
find	Utility that recursively searches directories for specified input
grep	Utility that searches files and directories for specified input
halt	Utility that stops the server
httpsdctl	Handles secure Apache Web processes
hwclock	Utility that allows access to the hardware clock
ifconfig	Indicates the network configuration executable. Ensures that the MAC address stays constant
iptables	Handles access restriction based on changes made to the Access Configuration page. See Configuring the Access List for Your Appliance, page 43-3 for more information about access configuration.
iptables-restore	Handles iptables file restoration
iptables-save	Handles saved changes to the iptables
kill	Utility that can be used to end a session and process
killall	Utility that can be used to end all sessions and processes
ksh	Public domain version of the Korn shell
logger	Utility that provides a way to access the syslog daemon from the command line
md5sum	Utility that prints checksums and block counts for specified files
mv	Utility that moves (renames) files
myisamchk	Indicates database table checking and repairing
mysql	Indicates a database process; multiple instances may appear
openssl	Indicates authentication certificate creation
perl	Indicates a perl process
ps	Utility that writes process information to standard output
sed	Utility used to edit one or more text files
sh	Public domain version of the Korn shell
shutdown	Utility that shuts down the appliance

Table 47-4 System Executables and Utilities (continued)

Executable	Description
sleep	Utility that suspends a process for a specified number of seconds
smtpclient	Mail client that handles email transmission when email event notification functionality is enabled
snmptrap	Forwards SNMP trap data to the SNMP trap server specified when SNMP notification functionality is enabled
snort (requires Protection)	Indicates that Snort is running
ssh	Indicates a Secure Shell (SSH) connection to the appliance
sudo	Indicates a sudo process, which allows users other than admin to run executables
top	Utility that displays information about the top CPU processes
touch	Utility that can be used to change the access and modification times of specified files
vim	Utility used to edit text files
wc	Utility that performs line, word, and byte counts on specified files



Using Backup and Restore

Backup and restoration is an essential part of any system maintenance plan. While each organization's backup plan is highly individualized, the ASA FirePOWER module provides a mechanism for archiving data so that data can be restored in case of disaster.

Note the following limitations about backup and restore:

- Backups are valid only for the product version on which you create them.
- You can restore a backup only when running the same version of the ASA FirePOWER module software as that used to create the backup.

Caution Do not use the backup and restore process to copy the configuration files between ASA FirePOWER modules. The configuration files include information that uniquely identifies an ASA FirePOWER module and cannot be shared.

Caution If you applied any intrusion rule updates, those updates are not backed up. You need to apply the latest rule update **after** you restore.

You can save backup files to the appliance or to your local computer.

See the following sections for more information:

- See Creating Backup Files, page 48-1 for information about creating backup files.
- See Creating Backup Profiles, page 48-3 for information about creating backup profiles that you can use later as templates for creating backups.
- See Uploading Backups from a Local Host, page 48-4 for information about uploading backup files from a local host.
- See Restoring the Appliance from a Backup File, page 48-4 for information about how to restore a backup file to the appliance.

Creating Backup Files

License: Any

You can perform backups of the ASA FirePOWER module using the module interface. To view and use existing system backups, go to the Backup Management page. You should periodically save a backup file that contains all of the configuration files required to restore the appliance, in addition to event data.

You may also want to back up the system when testing configuration changes so that you can revert to a saved configuration if needed. You can choose to save the backup file on the appliance or on your local computer.

You cannot create a backup file if your appliance does not have enough disk space; backups may fail if the backup process uses more than 90% of available disk space. If necessary, delete old backup files, transfer old backup files off the appliance.

As an alternative, or if your backup file is larger than 4GB, copy it via SCP to a remote host. Uploading a backup from your local computer does not work on backup files larger than 4GB.

Caution If you configured any interface associations with security zones, these associations are not backed up. You must reconfigure them after you restore. For more information, see Working with Security Zones, page 2-32.

To create a backup file of the ASA FirePOWER module:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Tools > Backup/Restore**.

The Backup Management page appears.

Step 2 Click **Device Backup**.

The Create Backup page appears.

Step 3 In the **Name** field, type a name for the backup file. You can use alphanumeric characters, punctuation, and spaces.

Step 4 Optionally, to be notified when the backup is complete, select the **Email** check box and type your email address in the accompanying text box.

Note To receive email notifications, you must configure a relay host as described in Configuring a Mail Relay Host and Notification Address, page 43-6.

Step 5 Optionally, to use secure copy (`SCP`) to copy the backup archive to a different machine, select the **Copy when complete** check box, then type the following information in the accompanying text boxes:

- in the **Host** field, the hostname or IP address of the machine where you want to copy the backup
- in the **Path** field, the path to the directory where you want to copy the backup
- in the **User** field, the user name you want to use to log into the remote machine
- in the **Password** field, the password for that user name
  If you prefer to access your remote machine with an SSH public key instead of a password, you must copy the contents of the **SSH Public Key** field to the specified user's `authorized_keys` file on that machine.

With this option cleared, the system stores temporary files used during the backup on the remote server; temporary files are **not** stored on the remote server when this option is selected.

Tip Cisco recommends that you periodically save backups to a remote location so the appliance can be restored in case of system failure.

Step 6 You have the following options:

- To save the backup file to the appliance, click **Start Backup**.

The backup file is saved in the `/var/sf/backup` directory.

When the backup process is complete, you can view the file on the Restoration Database page. For information about restoring a backup file, see Restoring the Appliance from a Backup File, page 48-4.

- To save this configuration as a backup profile that you can use later, click **Save As New**.

  You can modify or delete the backup profile by selecting **Configuration > ASA FirePOWER Configuration > Tools > Backup/Restore**, then clicking **Backup Profiles**. See Creating Backup Profiles, page 48-3 for more information.

Creating Backup Profiles

License: Any

You can use the Backup Profiles page to create backup profiles that contain the settings that you want to use for different types of backups. You can later select one of these profiles when you back up the files on your appliance.

Tip When you create a backup file as described in Creating Backup Files, page 48-1, a backup profile is automatically created.

To create a backup profile:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Tools > Backup/Restore**.

The Backup Management page appears.

Step 2 Click the **Backup Profiles** tab.

The Backup Profiles page appears with a list of existing backup profiles.

Tip You can click the edit icon ( ) to modify an existing profile or click the delete icon ( ) to delete a profile from the list.

Step 3 Click **Create Profile**.

The Create Backup page appears.

Step 4 Type a name for the backup profile. You can use alphanumeric characters, punctuation, and spaces.

Step 5 Configure the backup profile according to your needs.

See Creating Backup Files, page 48-1 for more information about the options on this page.

Step 6 Click **Save As New** to save the backup profile.

The Backup Profiles page appears and your new profile appears in the list.

Uploading Backups from a Local Host

License: Any

If you download a backup file to your local host using the download function described in the Backup Management table, you can upload it to an ASA FirePOWER module.

If your backup file contains PKI objects, private keys associated with internal CA and internal certificate objects are re-encrypted on upload with a randomly generated key.

Tip You cannot upload a backup larger than 4GB from your local host. As an alternative, copy the backup via SCP to a remote host and retrieve it from there.

To upload a backup from your local host:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Tools > Backup/Restore**.

The Backup Management page appears.

Step 2 Click **Upload Backup.**

The Upload Backup page appears.

Step 3 Click **Choose File** and navigate to the backup file you want to upload.

After you select the file to upload, click **Upload Backup**.

Step 4 Click **Backup Management** to return to the Backup Management page.

The backup file is uploaded and appears in the backup list. After the ASA FirePOWER moduleverifies the file integrity, refresh the Backup Management page to reveal detailed file system information.

Restoring the Appliance from a Backup File

License: Any

You can restore the appliance from backup files using the Backup Management page. To restore a backup, the VDB version in the backup file must match the current VDB version on your appliance. After you complete the restoration process, you **must** apply the latest Cisco Rule Update.

Caution Do not restore backups created on virtual Firepower Management Centers to physical Firepower Management Centers — this may stress system resources. If you must restore a virtual backup on a physical Firepower Management Center, contact Support.

If your backup file contains PKI objects, private keys associated with internal CA and internal certificate objects are reencrypted on upload with a randomly generated key.

If you use local storage, backup files are saved to `/var/sf/backup`, which is listed with the amount of disk space used in the `/var` partition at the bottom of the Backup Management page.

Note If you add licenses after a backup has completed, these licenses will not be removed or overwritten if this backup is restored. To prevent a conflict on restore, remove those licenses before restoring the backup, noting where the licenses were used, and add and reconfigure them after restoring the backup. If a conflict occurs, contact Support.

The following table describes each column and icon on the Backup Management page.

Table 48-1 Backup Management

Functionality	Description
System Information	The originating appliance name, type, and version. Note that you can only restore a backup to an identical appliance type and version.
Date Created	The date and time that the backup file was created
File Name	The full name of the backup file
VDB Version	The build of the vulnerability database (VDB) running on the appliance at the time of backup.
Location	The location of the backup file
Size (MB)	The size of the backup file, in megabytes
View	Click the name of the backup file to view a list of the files included in the compressed backup file.
Restore	Click with the backup file selected to restore it on the appliance. If your VDB version does not match the VDB version in the backup file, this option is disabled.
Download	Click with the backup file selected to save it to your local computer.
Delete	Click with the backup file selected to delete it.
Move	When you have a previously created local backup selected, click to send the backup to the designated remote backup location.

To restore the appliance from a backup file:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Tools > Backup/Restore**.

The Backup Management page appears.

Step 2 To view the contents of a backup file, click the name of the file.

The manifest appears, listing the name of each file, its owner and permissions, and its file size and date.

Step 3 Click **Backup Management** to return to the Backup Management page.

Step 4 Select the backup file that you want to restore and click **Restore**.

The Restore Backup page appears.

Note that if the VDB version in the backup does not match the VDB version currently installed on your appliance, the **Restore** button is grayed out.

Caution This procedure overwrites all configuration files.

Step 5 To restore files, select **Replace Configuration Data.**

Step 6 Click **Restore** to begin the restoration.

The appliance is restored using the backup file you specified.

Step 7 Reboot the appliance.

Step 8 Apply the latest Cisco Rule Update to reapply rule updates.

Step 9 Redeploy policies to the restored system.



Generating Troubleshooting Files

In some cases, if you have a problem with your appliance, Support may ask you to generate troubleshooting files to help them diagnose the problem. You can select any of the options listed in the following table to customize the troubleshooting data that the ASA FirePOWER module reports.

Table A-1 Selectable Troubleshoot Options

This option...	Reports...
Snort Performance and Configuration	data and configuration settings related to Snort on the appliance
Hardware Performance and Logs	data and logs related to the performance of the appliance hardware
System Configuration, Policy, and Logs	configuration settings, data, and logs related to the current system configuration of the appliance
Detection Configuration, Policy, and Logs	configuration settings, data, and logs related to detection on the appliance
Interface and Network Related Data	configuration settings, data, and logs related to inline sets and network configuration of the appliance
Discovery, Awareness, VDB Data, and Logs	configuration settings, data, and logs related to the current discovery and awareness configuration on the appliance
Upgrade Data and Logs	data and logs related to prior upgrades of the appliance
All Database Data	all database-related data that is included in a troubleshoot report
All Log Data	all logs collected by the appliance database
Network Map Information	current network topology data

Note that some options overlap in terms of the data they report, but the troubleshooting files will not contain redundant copies, regardless of what options you select.

For more information, see the following sections:

- Generating Appliance Troubleshooting Files, page A-1
- Downloading Troubleshooting Files, page A-2

Generating Appliance Troubleshooting Files

License: Any

Use the following procedure to generate customized troubleshooting files that you can send to Support.

To generate troubleshooting files:

Step 1 In ASDM, select **Configuration > ASA FirePOWER Configuration > Tools > Troubleshooting**.

Step 2 Click **Generate Troubleshooting Files**.

The Troubleshooting Options pop-up window appears.

Step 3 Select **All Data** to generate all possible troubleshooting data, or select individual check boxes to customize your report. For more information, see the Selectable Troubleshoot Options table.

Step 4 Click **OK**.

The ASA FirePOWER module generates the troubleshooting files. You can monitor the file generation process in the task queue (**Monitoring > ASA FirePOWER Monitoring > Task Status**).

Step 5 Continue with the procedure in the next section, Downloading Troubleshooting Files.

Downloading Troubleshooting Files

License: Any

Use the following procedure to download copies of your generated troubleshooting files.

To download troubleshooting files:

Step 1 In ASDM, select **Monitoring > ASA FirePOWER Monitoring > Task Status**.

The Task Status page appears.

Step 2 Find the task that corresponds to the troubleshooting files you generated.

Step 3 After the appliance generates the troubleshooting files and the task status changes to `Completed`, click **Click to retrieve generated files**.

Step 4 Follow your browser's prompts to download the files.

The files are downloaded in a single `.tar.gz` file.

Step 5 Follow the directions from Support to send the troubleshooting files to Cisco.



Importing and Exporting Configurations

You can use the Import/Export feature to copy several types of configurations, including policies, from one appliance to another appliance of the same type. Configuration import and export is not intended as a backup tool, but can be used to simplify the process of adding new ASA FirePOWER modules.

You can import and export the following configurations:

- access control policies and their associated network analysis, SSL, and file policies
- intrusion policies
- system policies
- alert responses

To import an exported configuration, both ASA FirePOWER modules must be running the same software version. To import an exported intrusion or access control policy, the rule update versions on both appliances must also match.

For more information, see the following sections:

- Exporting Configurations, page B-1
- Importing Configurations, page B-3

Exporting Configurations

License: Any

You can export a single configuration, or you can export a set of configurations (of the same type or of different types) at once. When you later import the package onto another appliance, you can choose which configurations in the package to import.

When you export a configuration, the appliance also exports revision information for that configuration. The ASA FirePOWER module uses that information to determine whether you can import that configuration onto another appliance; you cannot import a configuration revision that already exists on an appliance.

In addition, when you export a configuration, the appliance also exports system configurations that the configuration depends on.

Tip Many list pages in the ASA FirePOWER module include an export icon ( ) next to list items. Where this icon is present, you can use it as a quick alternative to the export procedure that follows.

You can export the following configurations:

- *Alert responses* — An alert response is a set of configurations that allows the ASA FirePOWER module to interact with the external system where you plan to send the alert.
- *Access control policies* — Access control policies include a variety of components that you can configure to determine how the system manages traffic on your network. These components include access control rules; associated intrusion, file, and network analysis, and SSL policies; and objects the rules and policies use, including intrusion variable sets. Exporting an access control policy exports all settings and components for the policy except (where present) URL reputations and categories, which are equivalent across appliances and which users cannot change. Note that to import an access control policy, the rule update version on the exporting and importing ASA FirePOWER module must match.

  If an access control policy that you export, or the SSL policy it invokes, contains rules that reference geolocation data, the importing module's geolocation database (GeoDB) update version is used.
- *Intrusion policies* — Intrusion policies include a variety of components that you can configure to inspect your network traffic for intrusions and policy violations. These components are intrusion rules that inspect the protocol header values, payload content, and certain packet size characteristics, and other advanced settings.

  Exporting an intrusion policy exports all settings for the policy. For example, if you choose to set a rule to generate events, or if you set SNMP alerting for a rule, or if you turn on the sensitive data preprocessor in a policy, those settings remain in place in the exported policy. Custom rules, custom rule classifications, and user-defined variables are also exported with the policy.

  Note that if you export an intrusion policy that uses a layer that is shared by a second intrusion policy, that shared layer is copied into the policy you are exporting and the sharing relationship is broken. When you import the intrusion policy on another appliance, you can edit the imported policy to suit your needs, including deleting, adding, and sharing layers.

  If you export an intrusion policy from one ASA FirePOWER module to another, the imported policy may behave differently if the second ASA FirePOWER module has differently configured default variables.

  Note You cannot use the Import/Export feature to update rules created by the Vulnerability Research Team (VRT). Instead, download and apply the latest rule update version; see Importing Rule Updates and Local Rule Files, page 46-9.

- *System policies* — A system policy controls the aspects of an ASA FirePOWER module that are likely to be similar to other ASA FirePOWER modules in your deployment, including time settings, SNMP settings, and so on.

Note Depending on the number of configurations being exported and the number of objects those configurations reference, the export process may take several minutes.

To export one or more configurations:

Step 1 Make sure that the ASA FirePOWER module where you are exporting the configurations and the ASA FirePOWER module where you plan to import the configurations are running the same version. If you are exporting an intrusion or access control policy, make sure that the rule update versions match.

If the versions of the ASA FirePOWER module (and, if applicable, the rule update versions) do not match, the import will fail.

Step 2 Select **Configuration > ASA FirePOWER Configuration > Tools > Import Export**.

The Import/Export page appears, including a list of the configurations on the ASA FirePOWER module. Note that configuration categories with no configurations to export do not appear in this list.

Tip You can click the collapse icon ( ) next to a configuration type to collapse the list of configurations. Click the expand folder icon ( ) next to an configuration type to reveal configurations.

Step 3 Select the check boxes next to the configurations you want to export and click **Export**.

Step 4 Follow the prompts to save the exported package to your computer.

Importing Configurations

License: Any

After you export a configuration from an ASA FirePOWER module, you can import it onto a different module as long as that module supports it.

Depending on the type of configuration you are importing, you should keep the following points in mind:

- You must make sure that the ASA FirePOWER module where you import a configuration is running the same version as the ASA FirePOWER module you used to export the configuration. If you are importing an intrusion or access control policy, the rule update versions on both appliances must also match. If the versions do not match, the import will fail.
- If you import an access control policy that evaluates traffic based on zones, you must map the zones in the imported policy to zones on the importing ASA FirePOWER module. When you map zones, their types must match. Therefore, you must create any zone types you need on the importing ASA FirePOWER module before you begin the import. For more information about security zones, see Working with Security Zones, page 2-32.
- If you import an access control policy that includes an object or object group that has an identical name to an existing object or group, you must rename the object or group.
- If you import an access control policy or an intrusion policy, the import process replaces existing default variables in the default variable set with the imported default variables. If your existing default variable set contains a custom variable not present in the imported default variable set, the unique variable is preserved.
- If you import an intrusion policy that used a shared layer from a second intrusion policy, the export process breaks the sharing relationship and the previously shared layer is copied into the package. In other words, imported intrusion policies do not contain shared layers.

Note You cannot use the Import/Export feature to update rules created by the Vulnerability Research Team (VRT). Instead, download and apply the latest rule update version; see Importing Rule Updates and Local Rule Files, page 46-9.

Because you can export several configurations in a single package, when you import the package you must choose which configurations in the package to import.

When you attempt to import a configuration, your ASA FirePOWER module determines whether that configuration already exists on the appliance. If a conflict exists, you can:

- keep the existing configuration,
- replace the existing configuration with a new configuration,
- keep the newest configuration, or
- import the configuration as a new configuration.

If you import a configuration and then later make a modification to the configuration on the destination system, and then re-import the configuration, you must choose which version of the configuration to keep.

Depending on the number of configurations being imported and the number of objects those configurations reference, the import process may take several minutes.

To import one or more configurations:

Step 1 Make sure that the ASA FirePOWER module where you are exporting the configurations and the module where you plan to import the configurations are running the same version. If you want to import an intrusion or access control policy, you must also make sure that the rule update versions match.

If the versions of the ASA FirePOWER module (and, if applicable, the rule update versions) do not match, the import will fail.

Step 2 Export the configurations you want to import; see Exporting Configurations, page B-1.

Step 3 On the appliance where you want to import the configurations, select **Configuration > ASA FirePOWER Configuration > Tools > Import Export**.

The Import Export page appears.

Tip Click the collapse icon () next to a configuration type to collapse the list of configurations. Click the expand folder icon () next to a configuration type to reveal configurations.

Step 4 Click **Upload Package**.

The Upload Package page appears.

Step 5 You have two options:

- Type the path to the package you want to upload.
- Click **Upload File** to locate the package.

Step 6 Click **Upload**.

The result of the upload depends on the contents of the package:

- If the configurations in the package exactly match versions that already exist on your appliance, a message displays indicating that the versions already exist. The appliance has the most recent configurations, so you do not need to import them.
- If there is an ASA FirePOWER module or (if applicable) rule update version mismatch between your appliance and the appliance where the package was exported, a message appears, indicating that you cannot import the package. Update the ASA FirePOWER module or the rule update version and attempt the process again.
- If the package contains any configurations or rule versions that do not exist on your appliance, the Package Import page appears. Continue with the next step.

Step 7 Select the configurations you want to import and click **Import**.

The import process resolves, with the following results:

- If the configurations you import do not have previous revisions on your ASA FirePOWER module, the import completes automatically and a success message appears. Skip the rest of the procedure.
- If you are importing an access control policy that includes security zones, the Access Control Import Resolution page appears. Continue with step 8.
- If the configurations you import do have previous revisions on your appliance, the Import Resolution page appears. Continue with step 9.

Step 8 Next to each incoming security zone, select an existing local security zone of a matching type to map to and click **Import**.

Return to step 7.

Step 9 Expand each configuration and select the appropriate option:

- To keep the configuration on your appliance, select **Keep existing**.
- To replace the configuration on your appliance with the imported configuration, select **Replace existing**.
- To keep the newest configuration, select **Keep newest**.
- To save the imported configuration as a new configuration, select **Import as new** and, optionally, edit the configuration name.

  If you are importing an access control policy that includes a file policy with either the clean list or custom detection list enabled, the **Import as new** option is not available.
- If you are importing an access control policy or saved search that includes a dependent object, either accept the suggested name or rename the object. The system always imports these dependent objects as new. You do not have the option to keep or to replace existing objects. Note that the system treats objects and object groups in the same manner.

Step 10 Click **Import**.

The configurations are imported.



Viewing the Status of Long-Running Tasks

Some tasks that you can perform on the ASA FirePOWER module, such as applying a policy or installing updates, do not complete instantly and require some time to run. You can check the progress of these long-running tasks in the task queue. The task queue also reports when they are successfully or unsuccessfully resolved.

For more information, see the following sections:

- Viewing the Task Queue, page C-1
- Managing the Task Queue, page C-2

Viewing the Task Queue

License: Any

When you perform long-running tasks, such as applying a policy or installing updates, the status of these tasks is reported in the task queue. The task queue provides information about complex tasks and reports when they are complete.

You view the task queue on the Task Status page, which automatically refreshes every 10 seconds.

The Job Summary section displays the state of the tasks listed on the page, as described in the following table.

Table C-1 Task Queue Task Types

Task Type	Description
Running	The number of tasks currently in progress.
Waiting	The number of tasks waiting for an in-progress task to complete before running.
Completed	The number of tasks that completed successfully.
Retrying	The number of tasks that are automatically retrying. Note that not all tasks are permitted to try again.
Stopped	The number of tasks that were interrupted due to a system update. Stopped tasks cannot be resumed; you must manually delete them from the task queue.
Failed	The number of tasks that did not complete successfully.

The Jobs section provides information about each task, including a brief description, when the task was launched, the current status of the task, and when the status last changed. Tasks of the same type appear together in a task group.

To make sure that the Task Status page loads quickly, once per week, the ASA FirePOWER module removes from the queue all completed, failed, and stopped tasks that are over a month old, as well the oldest tasks from any task group that contains over 1000 tasks. You can also manually remove tasks from the queue; see Managing the Task Queue for directions.

To view the task queue:

Step 1 You have two options:

- If you manually launched the task, click the **Task Status** link in the notification box that appeared when you launched the task.

  The Task Status page appears in a pop-up window.
- If you scheduled a task, or if a task was launched from a page you are not viewing, select **Monitoring > ASA FirePOWER Monitoring > Task Status**.

  The Task Status page appears.

For information on the actions you can perform on the Task Status page, see Managing the Task Queue.

Managing the Task Queue

License: Any

There are several actions you can perform while viewing the task queue (see Viewing the Task Queue, page C-1), as described in the following table.

Table C-2 Task Queue Actions

To...	You can...
remove all completed tasks from the task queue	click **Remove Completed Jobs**.
remove all failed tasks from the task queue	click **Remove Failed Jobs**.
remove a single task from the task queue	click the delete icon ( ) next to the task you want to delete. Note that you cannot delete a running task. If you need to delete a running task (for example, if a task repeatedly fails), contact Support.
collapse a task group and hide tasks	click the open folder icon ( ) next to the expanded task group.
expand a task group and view tasks	click the closed folder icon ( ) next to the collapsed task group.



Security, Internet Access, and Communication Ports

To safeguard the ASA FirePOWER module, you should install it on a protected internal network. Although the ASA FirePOWER module is configured to have only the necessary services and ports available, you must make sure that attacks cannot reach it from outside the firewall.

Also note that specific features of the ASA FirePOWER module require an Internet connection. By default, the ASA FirePOWER module is configured to directly connect to the Internet. Additionally, the system requires certain ports remain open for secure appliance access and so that specific system features can access the local or Internet resources to operate correctly.

For more information, see:

- Internet Access Requirements, page D-1
- Communication Ports Requirements, page D-2

Internet Access Requirements

By default, the ASA FirePOWER module is configured to directly connect to the Internet on ports 443/tcp (HTTPS) and 80/tcp (HTTP), which are open by default on the ASA FirePOWER module; see Communication Ports Requirements, page D-2.

The following table describes the Internet access requirements of specific features of the ASA FirePOWER module.

Table D-1 ASA FirePOWER module Feature Internet Access Requirements

Feature	Internet access is required to...
intrusion rule, VDB, and GeoDB updates	download or schedule the download of a intrusion rule, GeoDB, or VDB update directly to an appliance.
network-based AMP	perform malware cloud lookups.
Security Intelligence filtering	download Security Intelligence feed data from an external source, including the Intelligence Feed.
system software updates	download or schedule the download of a system update directly to an appliance.

Table D-1 ASA FirePOWER module Feature Internet Access Requirements (continued)

Feature	Internet access is required to...
URL filtering	download cloud-based URL category and reputation data for access control, and perform lookups for uncategorized URLs.
whois	request whois information for an external host.

Communication Ports Requirements

Open ports allow:

- access to an appliance's user interface
- secure remote connections to an appliance
- certain features of the system to access the local or Internet resources they need to function correctly

In general, feature-related ports remain closed until you enable or configure the associated feature.

Caution Do **not** close an open port until you understand how this action will affect your deployment.

For example, closing port 25/tcp (SMTP) outbound on a manage device blocks the device from sending email notifications for individual intrusion events (see Configuring External Alerting for Intrusion Rules, page 39-1).

The following table lists the open ports required so that you can take full advantage of ASA FirePOWER module features.

Table D-2 Default Communication Ports for ASA FirePOWER module Features and Operations

Port	Description	Direction	Is Open to...
22/tcp	SSH/SSL	Bidirectional	allow a secure remote connection to the appliance.
25/tcp	SMTP	Outbound	send email notices and alerts from the appliance.
53/tcp	DNS	Outbound	use DNS.
67/udp 68/udp	DHCP	Outbound	use DHCP. **Note** These ports are **closed** by default.
		Bidirectional	update custom and third-party Security Intelligence feeds via HTTP. download URL category and reputation data (port 443 also required).
161/udp	SNMP	Bidirectional	allow access to an appliance's MIBs via SNMP polling.
162/udp	SNMP	Outbound	send SNMP alerts to a remote trap server.
389/tcp 636/tcp	LDAP	Outbound	communicate with an LDAP server for external authentication.
389/tcp 636/tcp	LDAP	Outbound	obtain metadata for detected LDAP users.

Table D-2 Default Communication Ports for ASA FirePOWER module Features and Operations (continued)

Port	Description	Direction	Is Open to...
443/tcp	HTTPS	Inbound	access an appliance's user interface.
443/tcp	HTTPS cloud comms.	Bidirectional	obtain: • software, intrusion rule, VDB, and GeoDB updates • URL category and reputation data (port 80 also required) • the Intelligence Feed and other secure Security Intelligence feeds • malware dispositions for files detected in network traffic
			download software updates using the device's local user interface.
514/udp	syslog	Outbound	send alerts to a remote syslog server.
8305/tcp	appliance comms.	Bidirectional	securely communicate between appliances in a deployment. **Required.**
8307/tcp	host input client	Bidirectional	communicate with a host input client.